

---

# **The metaphysics of enactivism: Realism and the egocentric profile**

*Geoff Allen*

6867138

---

*Master thesis (History and philosophy of science)*

*Friday 23 July 2021*

*First supervisor: Dr Brandt van der Gaast*

*Second supervisor: Dr Janneke van Lith*

*Words: 36,080*

*(excl. bibliography and footnotes)*

## Table of Contents

<b>1. General introduction</b>	6
1.1. The big picture	6
1.2. In the spirit of symmetry	9
1.3. Enactivism, egocentricity and realism	10
 <b>Block A: Enactivism and representationalism</b>	 15
<b>2. Enactivism: A taxonomy</b>	15
2.1. Masters of movement	15
2.2. Positions in the representationalism debate	18
2.2.1. Cognitivism (pro-representationalism)	18
2.2.2. Enactivism (anti-representationalism)	19
2.2.3. Sidebar: How does Tiger Woods perceive the world?	21
2.2.4. A spectrum of positions	24
2.3. What is representation?	26
2.3.1. The function of carrying information	26
2.3.2. A continuum of representing	28
 <b>3. Enactivism: Assessed</b>	 30
3.1. Assessment overview	30
3.2. Gibson: True anti-representationalism	31
3.3. Varela, Thompson and Rosch: Anti-foundationalist enactivism	34
3.4. O'Regan and Noë: Sensorimotor enactivism	37
3.4.1. Sensorimotor contingencies	37
3.4.2. Action, exploration and the world as an outside memory	40
3.4.3. Internal processing and sensorimotor knowledge	42
3.4.4. Representing expectations	45
3.4.5. Representing an index	48

3.4.6.	<i>Rehabilitating enactivism</i> .....	49
3.5.	<i>Clark: Minimal representing</i> .....	51
3.5.1.	<i>An ecumenical approach to representationalism</i> .....	51
3.5.2.	<i>Minimal versus maximal representations</i> .....	52
4.	<b><i>Enactivism: Cognitive science and empirical evidence</i></b> .....	55
4.1.	<i>Evidence for enactivism</i> .....	55
4.1.1.	<i>Perception does not involve detailed internal representations</i> .....	55
4.1.2.	<i>Action guides perception</i> .....	56
4.2.	<i>Evidence against enactivism</i> .....	58
4.2.1.	<i>The significance of sensorimotor contingencies</i> .....	58
4.3.	<i>Closing comments on representationalism</i> .....	59
	<b><i>Block B: Enactivism and egocentricity</i></b> .....	62
5.	<b><i>Piercing the soap bubble</i></b> .....	62
5.1.	<i>What is egocentricity?</i> .....	62
5.1.1.	<i>From enactivism to egocentricity</i> .....	62
5.1.2.	<i>The world inside a soap bubble</i> .....	64
5.1.3.	<i>Piercing the soap bubble</i> .....	65
5.2.	<i>Egocentricity and allocentricity in spatial navigation</i> .....	68
5.2.1.	<i>Defining egocentric and allocentric</i> .....	68
5.2.2.	<i>Linking allocentricity and mind-independence</i> .....	71
5.2.3.	<i>Egocentricity: A world with me at the centre</i> .....	72
5.2.4.	<i>Allocentricity: Abstracting away from myself</i> .....	73
5.2.5.	<i>Elementary versus complex spatial representations</i> .....	76
5.2.6.	<i>Variation in human spatial abilities</i> .....	79
5.3.	<i>Applying egocentricity and allocentricity to enactivism</i> .....	83
5.3.1.	<i>Von Uexküll: The Umwelt</i> .....	83
5.3.2.	<i>Gibson: Ecological psychology and affordances</i> .....	86
5.3.3.	<i>Noë: Sensorimotor enactivism</i> .....	90

5.4.	<i>Perspective and object recognition</i> .....	93
5.4.1.	<i>The dual aspect of perceptual content</i> .....	93
5.4.2.	<i>Perspectival properties</i> .....	95
5.4.3.	<i>The debate about perspectival shape</i> .....	99
5.4.4.	<i>Experiments and findings</i> .....	102
5.5.	<i>Closing comments on egocentricity and allocentricity</i> .....	104
5.5.1.	<i>The dual nature of human perception</i> .....	104
5.5.2.	<i>The prospect of piercing the soap bubble</i> .....	106

## ***Block C: Enactivism and metaphysics* .....**

6.	<b><i>The matrix: Representationalism and realism</i></b> .....	109
6.1.	<i>Where representationalism and realism meet</i> .....	109
6.1.1.	<i>Entering the realism debate</i> .....	109
6.1.2.	<i>Defining ‘realism’ and ‘anti-realism’</i> .....	111
6.2.	<i>The matrix</i> .....	114
6.2.1.	<i>Connecting representationalism and realism</i> .....	114
6.2.2.	<i>The explicit and the tacit</i> .....	118
6.3.	<i>Plotting the points</i> .....	120
6.3.1.	<i>Gibson</i> .....	120
6.3.2.	<i>O’Regan and Noë</i> .....	123
6.3.3.	<i>Varela, Thompson and Rosch</i> .....	128
6.3.4.	<i>Von Uexküll</i> .....	132
6.3.5.	<i>Hutto and Myin</i> .....	134
6.3.6.	<i>Clark</i> .....	137
7.	<b><i>Enactivism and realism: Assessed</i></b> .....	141
7.1.	<i>Assessing the realism debate</i> .....	141
7.1.1.	<i>In the spirit of symmetrical realism</i> .....	141
7.1.2.	<i>The egocentric profile</i> .....	143
7.2.	<i>Embodiment: The argument from biology</i> .....	145

7.2.1.	<i>Egocentricity and anti-realism</i> .....	145
7.2.2.	<i>Egocentricity and realism</i> .....	147
7.2.3.	<i>Circularity: Begging the question of realism</i> .....	151
7.2.4.	<i>Biological egocentricity: From the species to the individual</i> .....	153
7.3.	<i>Embeddedness: The argument from geometry</i> .....	155
7.3.1.	<i>Embeddedness and vantage points</i> .....	155
7.3.2.	<i>Embeddedness: Externalism and extensive cognition</i> .....	158
7.3.3.	<i>Embeddedness and anti-realism</i> .....	160
7.3.4.	<i>Embeddedness and realism</i> .....	162
7.4.	<i>Closing comments on metaphysical realism</i> .....	166
8.	<b><i>General conclusion</i></b> .....	169
8.1.	<i>The metaphysics of enactivism</i> .....	169
8.1.1.	<i>Enactivism</i> .....	169
8.1.2.	<i>Egocentricity</i> .....	170
8.1.3.	<i>Metaphysics</i> .....	172
8.1.4.	<i>Realism and the egocentric profile</i> .....	174
9.	<b><i>Bibliography</i></b> .....	175

# 1. General introduction

## 1.1. The big picture


In this thesis project, I ask questions about human perception and metaphysical reality, taking *enactivism* as a case study. Enactivism is a research programme within philosophy of perception that makes certain claims about *representations*, *action*, *embodiment*, *embeddedness* and *externalism*. Over the course of this thesis project, I will discuss each of these elements. Crucially, enactivism entails a metaphysical position about the nature of perception: that perception arises from a dynamical interaction between organism and environment, grounded in the physical body of the organism. Enactivism is sometimes seen as radical, since it rejects the idea that perception involves inner representations of the outer environment.

This thesis project investigates *The metaphysics of enactivism*, which is one narrow slice of the enactivist literature. There are enactive approaches to a great number of topics, including ‘perception, intentionality, emotion, memory, social cognition and consciousness’.<sup>1</sup> This project cannot hope to provide a comprehensive assessment of the enactivist doctrine as a whole. Many stones will be left unturned. I will actively avoid questions that relate to consciousness, intentionality or epistemology, which are matters of importance for enactivism but are beyond the scope of this project. Instead, throughout this project, I will constrain myself to the ultimate question of metaphysics.

---

<sup>1</sup> Hutto and Myin, *Radicalizing enactivism*, 2.

How does this topic, the metaphysics of enactivism, fit into ‘the big picture’? It is useful to locate the topic against a much larger backdrop of connected subfields of philosophy. *Figure 1* shows the metaphysics of enactivism at the centre of a Venn diagram, at the centre of three interlocking spheres of research. First and foremost, enactivism is connected to the philosophy of mind; enactivists ask questions – and make claims – about the nature of cognition. Secondly, enactivism is connected to metaphysics, since enactivists interrogate the metaphysical arrangement between organism and environment, between perceiver and world. These are primarily questions about philosophy of perception, which I depict in the overlap between philosophy of mind and metaphysics, at *Figure 1*.



*Figure 1: The metaphysics of enactivism and the big picture*

There is a third and final subfield to note: the sphere of empirically motivated philosophy. Where this sphere intersects with philosophy of mind and metaphysics, researchers pursue ‘naturalised philosophy of mind’ and ‘naturalised metaphysics’, respectively. I take a firm metaphilosophical stance in favour of empirically motivated philosophy. I believe that such links to empirical work are crucial, especially when it comes to philosophy of mind.<sup>2</sup> Fortunately, many contemporary philosophers of mind draw heavily on neuroscience and cognitive science, and it has become a relatively mainstream approach.

Enactivism is a research programme with strong links to empirical cognitive science. Many of its claims are amenable to empirical testing and enactivists routinely refer to empirical research to support their arguments. In my view, this makes enactivism a highly promising research programme, and one which appeals to my motivations. This was an influential factor in choosing enactivism as a case study for this thesis project. In my survey of the metaphysics of enactivism, I will look to evidence from cognitive science. Thus, in this thesis project, I will be doing naturalised philosophy of mind, and – to a certain degree – naturalised metaphysics. That is the big picture. In this thesis project, I will be working at the intersection of metaphysics, philosophy of mind and empirically motivated philosophy.

---

<sup>2</sup> My metaphilosophical stance, in favour of making links to empirical work, is beyond the scope of this project. It would require a second thesis project, of a similar length, to develop and defend such a position. I address it briefly, in this introduction, since it is part of ‘the big picture’ – studying perception and metaphysics in an empirical context.



## 1.2. In the spirit of symmetry

In this thesis project, I'm motivated by a spirit of symmetry – accounting for both mind and world, both perceiver and object. The subtitle of my project, *Realism and the egocentric profile*, highlights these two elements: on the one hand, metaphysical reality, outside of any subject's perspective; on the other hand, the unique 'profile' of the perceiver.

I borrow the term 'symmetry' from Brian Cantwell Smith. He calls his position 'symmetrical realism', since it involves deference to human activity *and* deference to the world. Any robust account must explain both the world and the place of the perceiver within it. On the one hand, 'We are here', and on the other hand, 'We are not the only things here'.<sup>3</sup> In my thesis project, I pursue symmetry. I wish to grant due deference to the place of the subject in perception, without undermining the persistence, invariance or regularity of the world.

My symmetrical approach comes with a pair of technical terms: *egocentricity* and *allocentricity*. I wish to affirm the reality of metaphysical invariances outside of human thought, that is to say: 'We are not the only things here'. I will use the label allocentricity or allocentric properties to describe these mind-independent invariances. At the same time, I wish to account for the perceiver, the agent, the human being. 'We are here,' the perceivers attest. Indeed, I will argue that *the egocentric profile* of the perceiver can never be erased from perception. Perception is constrained by the 'profile' of the perceiver: the inventory of facts about the perceiver and their place in

---

<sup>3</sup> Smith, *On the origin of objects*, 88.

the environment. These egocentric facts about the perceiver – from physical vantage point to metabolism and sensorimotor system – feature indelibly in perception. In this thesis project, I will account for the egocentric profile, and argue that we can still, nonetheless, account for allocentricity and metaphysical realism.

I investigate this symmetry in the context of enactivism. How do enactivists account for, on the one hand, egocentric facts about the perceiver and, on the other hand, the independence of the world? Enactivism provides an astute characterisation of the role of the perceiver, in terms of action, and in terms of the broader embodied and embedded framework. This characterisation of the role of the perceiver will be central to my account of the egocentric profile. We stand to learn lots from the enactivist treatment of egocentricity. When it comes to the independence of the world, there is widespread disagreement between enactivists. Some enactivists take an anti-realist stance, arguing that the dynamical role of the perceiver in perception precludes the independence of the world. Other enactivists take a realist stance, in favour of the independence of the world. Ultimately, I will provide an explicit defence of realism in the context of enactivism; I will demonstrate that we can reconcile the notion of the egocentric profile with metaphysical realism. That, hopefully, will be my accomplishment of symmetry.

### **1.3. Enactivism, egocentricity and realism**

In the final part of this introduction, I will explain the structure of the thesis. My thesis project revolves around three core elements – enactivism, egocentricity and realism

– all of which I have briefly introduced. These three elements come together to form my central research question:

**The metaphysics of enactivism: Realism and the egocentric profile**

Research question

*In the context of **enactivism** and **the egocentric profile**,  
what are the implications (or what is the evidence) for metaphysical **realism**?*

In this thesis project, I aim to synthesise these three elements – enactivism, egocentricity and realism. In *Table 1*, I provide an overview of the ‘pairwise’ interactions between these elements. In §1.2, I discussed the context of symmetry, which involves the interaction between egocentricity (the place of the perceiver) and realism (the world). When discussing egocentricity and realism together, we can ask such questions about symmetry between subject and world. My ultimate claim – in favour of reconciling egocentricity and reality – will be presented in §7.

There is a related interaction between enactivism and realism. Some enactivists argue that enactivism entails anti-realism, since the world is ‘enacted’ through the perceptual process and is therefore not independent of the subject. Conversely, other enactivists hold the world to be independent of perceiving subjects. Assessing these different positions on realism will be a central focus in §6.

There is also an interaction between enactivism and egocentricity, which I discuss primarily in §5. I argue that the central claims of enactivism should be understood in terms of egocentricity. In a nutshell, enactivists appear to be making a claim about the egocentric character of perception. They claim that perception can only be understood in the context of the activity of the organism, arising from action, embodiment (the physical body) and embeddedness (the physical environment). So, for enactivists, perception has the character of being intrinsically connected to egocentric facts about the perceiver. Each of these pairwise interactions will be made clear in the relevant section.

Element		Element		Interaction	§
Enactivism	+	Egocentricity	=	Action, embodiment, embeddedness	5
Enactivism	+	Realism	=	World as independent or enacted?	6
Egocentricity	+	Realism	=	Symmetry? Reconcilable?	7

*Table 1: Pairwise interactions between enactivism, egocentricity and realism*

I approach my main research question in three main blocks. The three main elements – enactivism, egocentricity and realism – make up these three blocks: Block A, Block B and Block C. Each of these blocks investigates a sub-question: Sub-question A, Sub-question B and Sub-question C. An overview box is included at the beginning and end of each block. The three overview boxes appear as follows:

## Overview of Block A: Enactivism and representationalism

(§2, §3 and §4)

- **Sub-question A:** *To what extent has enactivism challenged or influenced representationalism?*
- **Conclusion A1:** Perception is guided by both action and minimal representations.
- **Conclusion A2:** Enactivism must be rehabilitated to include minimal representations.

## Overview of Block B: Enactivism and egocentricity

(§5)

- **Sub-question B:** *To what extent should human perception be understood as 'egocentric' or 'allocentric'?*
- **Conclusion B1:** Perception is guided by both egocentric and allocentric elements.
- **Conclusion B2:** Enactivism should be understood in terms of egocentricity and allocentricity.

## Overview of Block C: Enactivism and metaphysics

(§6, §7)

- **Sub-question C:** *What is the link between enactivism and realism, as claimed by other authors and as justified by the evidence?*
- **Conclusion C1:** Enactivists disagree about metaphysical implications, claiming realism, anti-realism or agnosticism.
- **Conclusion C2:** Enactivism (and, in particular, the egocentric profile of a human perceiver) is compatible with metaphysical realism.
  - **The argument from biology:** Metaphysical realism is compatible with egocentric facts about embodiment.
  - **The argument from geometry:** Metaphysical realism is compatible with egocentric facts about embeddedness.

# Block A: Enactivism and representationalism

## 2. Enactivism: A taxonomy

### 2.1. Masters of movement

In 2019, golfer Tiger Woods won the Masters Tournament, beating the three runners-up by just one stroke. During the tournament, Woods had taken one fewer swipe at the tiny white ball. This was his fifth time winning the Masters and, on this occasion, he took home \$2.07 million.

Golf is an unusual sport in certain ways. Unlike football or basketball, golf does not use a standardised field. Every golf course is different. Each day, players encounter a new terrain, with its own slopes and vegetation, and – each day – the wind is different. Compared to other athletes, golfers face a somewhat unique challenge. In most other ball sports, players have an unlimited number of ‘strokes’. For example, in cricket, tennis, handball, basketball or football, players can strike the ball as many times as they like. In golf, *every strike counts*. This calls for an extremely high level of precision. Golfers are masters of observing and reacting; they are masters of *perception* and *action*.

On a warm April morning, Tiger Woods stands at the first tee, looking out over the fairway.<sup>4</sup> He *perceives* the course ahead. He judges the distance to the hole, the

---

<sup>4</sup> At the Masters Tournament, competitors play eighteen ‘holes’ per day. They begin each hole from a fixed location, known as a ‘tee’. Competitors must ‘tee off’, striking the ball down the main thoroughfare (or ‘fairway’) of that particular course.

curve of the trajectory and the direction of the wind. Then, he translates this visual information into *action*. Woods chooses the correct club and positions himself on the teeing ground. As he lowers his head, the peak of his cap casts a shadow across his face, shading his eyes from the dazzling sun. He takes one final glance at the fairway, takes a deep breath, and then strikes.

The question for the metaphysics of enactivism is: *how* does he perceive? How does Tiger Woods see the world? Enactivism attempts to answer this question. Enactivism is branch of cognitive science and philosophy of mind which says that *perception is grounded in action*. For the enactivist, there is a tight relationship between perception and action. Seeing is active process; it is an act of exploring the world. Vision is only possible through *mastery* of the laws of 'sensorimotor contingency'. An organism must be finely tuned to the patterns of physical interactions in the world around them. An organism must *master* these patterns in order to act and perceive.<sup>5</sup>

It is safe to conclude that Tiger Woods is a master of bodily movement. He has a stunning ability to hit a distant target, observing the dynamics of the world around him and making minute adjustments to his gestures. He has mastered the patterns ('sensorimotor contingencies') connecting perception and action.

However, the enactivist claim goes deeper. Enactivism makes a *positive* claim, about the connection between perception and action, as well as a *negative* claim, about the nature of representations. The negative claim goes as follows: in the brain of Tiger Woods, there is *no* representation (or *re*-presentation) of the golf course. When Tiger Woods 'sees' the course, his brain does not construct an internal model of how the

---

<sup>5</sup> O'Regan and Noë, 'A sensorimotor account of vision and visual consciousness', 943.



course looks, or what slopes he should take into account, or what rough grass he should avoid. Woods plays his game of golf without the use of representations. *Perception is unmediated*. As it happens, this negative claim is rather revolutionary.

In the following sections (§2, §3 and §4), I will evaluate the influence of enactivism. I will formulate an answer to **Sub-question A: *To what extent has enactivism challenged or influenced representationalism?*** The notion of a representation is the pivotal issue. Is perception *mediated* or *unmediated*? Do we perceive *by way of representations*, or do we perceive *directly*? These are the two opposing extremes of the debate: representationalism versus anti-representationalism.

I take a forensic approach, investigating enactivists' claims – both explicit and tacit – about perception and representation. In my survey, I cast a wide net. I aim to document the claims of the most influential enactivists: Alva Noë, J. Kevin O'Regan, Francisco Varela, Evan Thompson and Eleanor Rosch. I engage with some opponents of enactivism, like Jerry Fodor and Zenon Pylyshyn, and those who take more moderate positions, like Andy Clark. I also present the work of James J. Gibson, a historical figure, who has influenced the debate about perception.

In this section, I begin with 'Enactivism: A taxonomy' (§2). Then I engage more critically with the claims of enactivism – 'Enactivism: Assessed' (§3), and 'Enactivism: Cognitive science and empirical evidence' (§4).

## Overview of Block A: Enactivism and representationalism

(§2, §3 and §4)

- **Sub-question A:** *To what extent has enactivism challenged or influenced representationalism?*
- **Conclusion A1:** Perception is guided by both action and minimal representations.
- **Conclusion A2:** Enactivism must be rehabilitated to include minimal representations.

## 2.2. Positions in the representationalism debate

### 2.2.1. Cognitivism (pro-representationalism)

The representationalism debate is centred around the question: *is perception mediated or unmediated?* Cognitivists (pro-representationalists) argue that perception is mediated by representations. According to cognitivism, the organism's sensory receptors encode information about the environment, and then the mind builds up a detailed internal model.<sup>6</sup> Thus, the organism perceives by way of an internal model of the environment. The organism does not directly perceive the environment. Rather, perception is indirect – mediated by inner representations.

---

<sup>6</sup> Noë and Thompson, 'Introduction', 2.

Jerry Fodor and Zenon Pylyshyn are proponents of cognitivism. They defend the view that cognition involves internal acts of information processing, using such items as mental representations. Fodor and Pylyshyn's position in favour of representationalism also entails *computationalism*. There is a 'computational process' required to transform perceptual stimuli into representations and knowledge.<sup>7</sup> Thus, representationalism and computationalism are closely linked.

Cognitivism has been a mainstream position in philosophy of mind for several decades. Fodor and Pylyshyn describe cognitivism as the 'Establishment' view; they write in response to Gibson, a dissident who argues against Establishment theory.<sup>8</sup> Noë and Thompson describe cognitivism as the 'orthodoxy' or 'the orthodox view'.<sup>9</sup> Thus, representationalism has been a highly influential theory – seen by many as the mainstream conception of perception and cognition. This is where enactivism enters. Enactivists propose an alternative framework for perception and cognition. They set out in opposition to cognitivism, in opposition to the orthodoxy and the Establishment.

### **2.2.2. Enactivism (anti-representationalism)**

In opposition to the orthodoxy, there are several 'heterodox views'. These include: Gibson's theory of direct perception, part of his wider *ecological psychology*; Varela, Thompson and Rosch's enactivism; O'Regan and Noë's sensorimotor enactivism.<sup>10</sup> I

---

<sup>7</sup> Fodor and Pylyshyn, 'How direct is visual perception?', 167–8.

<sup>8</sup> Fodor and Pylyshyn, 'How direct is visual perception?', 167–8.

<sup>9</sup> Noë and Thompson, 'Introduction', 2–3.

<sup>10</sup> Noë and Thompson, 'Introduction', 3–6.

will analyse each of these views in detail, in §3. Gibson is seen as a historical influence for contemporary enactivist programmes, though enactivists distance themselves from certain aspects of his overall theory. All of these heterodox views are united in the belief that perception and action are inseparable; that is their motivation for rejecting the orthodoxy.<sup>11</sup>

Enactivists emphasise the role of action in perception, which is key to their rejection of inner representations. Charles Wallis and Wayne Wright label enactivism as an ‘active’ approach to studying vision. This is contrasted with ‘passive’ approaches – cognitivism, the orthodoxy or the Establishment view. One of the central tenets of passive vision research is the existence of detailed representations. By contrast, active approaches emphasise the role of action in vision, focusing on ‘motor planning and execution, memory, and attention’.<sup>12</sup> Enactivists reject internal models in perception, since perception is an operation arising from interaction with the environment, rather than an operation that involves internal processing. Enactivists reject the view that perception involves inner processing. Instead, perception is an *outer*, dynamical process. Perception arises from the interaction between the organism’s physical body and the environment.

Enactivism is closely linked to views about embodiment and embeddedness – claims that ‘cognition is “constituted” by bodily and environmental processes’.<sup>13</sup> According to the embodied view, cognition is constituted to some degree by bodily processes outside the brain, so the physical body of the organism (size, metabolism,

---

<sup>11</sup> Noë and Thompson, ‘Introduction’, 3.

<sup>12</sup> Wallis and Wright, ‘Enactivism’s vision’, 251.

<sup>13</sup> Newen, Gallagher, and De Bruin, ‘4E cognition’, 2–3.

sensorimotor system, etc.) shapes mentality. According to the embedded view, cognition is constituted by *extrabodily* processes in the environment of the organism.<sup>14</sup> Enactivists argue that perception arises from the interaction between organism and environment, between physical body and physical surroundings. Thus, cognition is embodied and embedded. For the enactivist, cognition is not constituted by the brain alone, but is shaped by properties of the body and the world. This is a distinctly *metaphysical* claim – key to the metaphysics of enactivism – as I will analyse in detail, in §7.

### **2.2.3. Sidebar: How does Tiger Woods perceive the world?**

So far, I have introduced cognitivism and enactivism. On the one hand, cognitivists argue that perception is mediated by representations – detailed internal models of the world – and that vision should be understood in a passive context. But what does this mean in real world terms? How does Tiger Woods perceive the world, according to the cognitivist?

---

<sup>14</sup> Newen, Gallagher, and De Bruin, '4E cognition', 5.



*Figure 2: The cognitivist's account of vision (mediated by representations)*

*Figure 2* is a crude depiction of the cognitivist's account of vision. The most crucial feature is an internal, action-neutral model of the world, represented in Tiger's brain. In the image, the small red box near Tiger's head indicates this re-presentation of the visual scene in his brain. For the cognitivist (like Fodor or Pylyshyn), vision is an act of internal information processing. Tiger's brain constructs a rich model of the environment, since (according to computationalism), perception is a computational process of translating visual stimuli into mental representations.

*Figure 3* shows the enactivist's contrasting account of vision. The enactivist argues that perception is unmediated, with no representations or internal models of the world. Instead, perception should be understood in the context of action and interaction with the environment. In *Figure 3*, there is no model of the world in Tiger's

brain. Instead, a red arrow shows a more direct engagement with environmental stimuli. When Tiger perceives the world, his brain does not reconstruct the visual scene, but rather interprets the contingencies linking perception and action. Each of Tiger's actions is accompanied by a contingent percept. *If I do this, then I will perceive that. If I swing the driver at this precise angle, then I will perceive the golf ball whizzing down the fairway.*<sup>15</sup> According to enactivism, Tiger is a masterful golfer precisely because he has mastered these contingencies between perceiving and action. He knows exactly which club and which angle and which minute movements to use, in order to produce the desired effect.

The enactivist's account is also linked to claims about embodiment and embeddedness. For the enactivist, perception arises from the interaction between physical body and physical environment. Perception (and cognition more generally) is constituted by bodily processes and extrabodily processes. In order to perceive, Tiger must explore the environment and interact dynamically. Thus, perception is an 'outer' activity of physical interaction, rather than an 'inner' activity of registration and processing.

---

<sup>15</sup> It might appear that this requires a model or representation of the contingencies – modelling if/then relationships. That would be an astute insight. As I will discuss in §3, O'Regan and Noë's sensorimotor enactivism depends on some, minimal representations.



Figure 3: The enactivist's account of vision (unmediated and grounded in action)

#### 2.2.4. A spectrum of positions

Thus far, the juxtaposition of cognitivism (orthodoxy) versus enactivism (heterodoxy) has been rather crude and lacking in nuance. We can now start to consider a wider spectrum of positions, from a wider range of scholarly fields. In *The embodied mind*, Francisco Varela, Evan Thompson and Eleanor Rosch present a visual map of the whole of cognitive science, as they see it.

Figure 4, a 'polar map', is taken from page 7 of *The embodied mind*.<sup>16</sup> Along one dimension, it can be read like a clock, identifying scholars from the related fields of artificial intelligence, linguistics, philosophy, cognitive psychology and neuroscience. Along its other dimension, the diagram shows three concentric rings: (i) cognitivism, (ii) emergence and (iii) enactivism. These are three competing research programmes

---

<sup>16</sup> Varela, Thompson and Rosch, *The embodied mind*, 7.



in cognitive science, each spanning several fields, and each with its own fundamental assumptions.

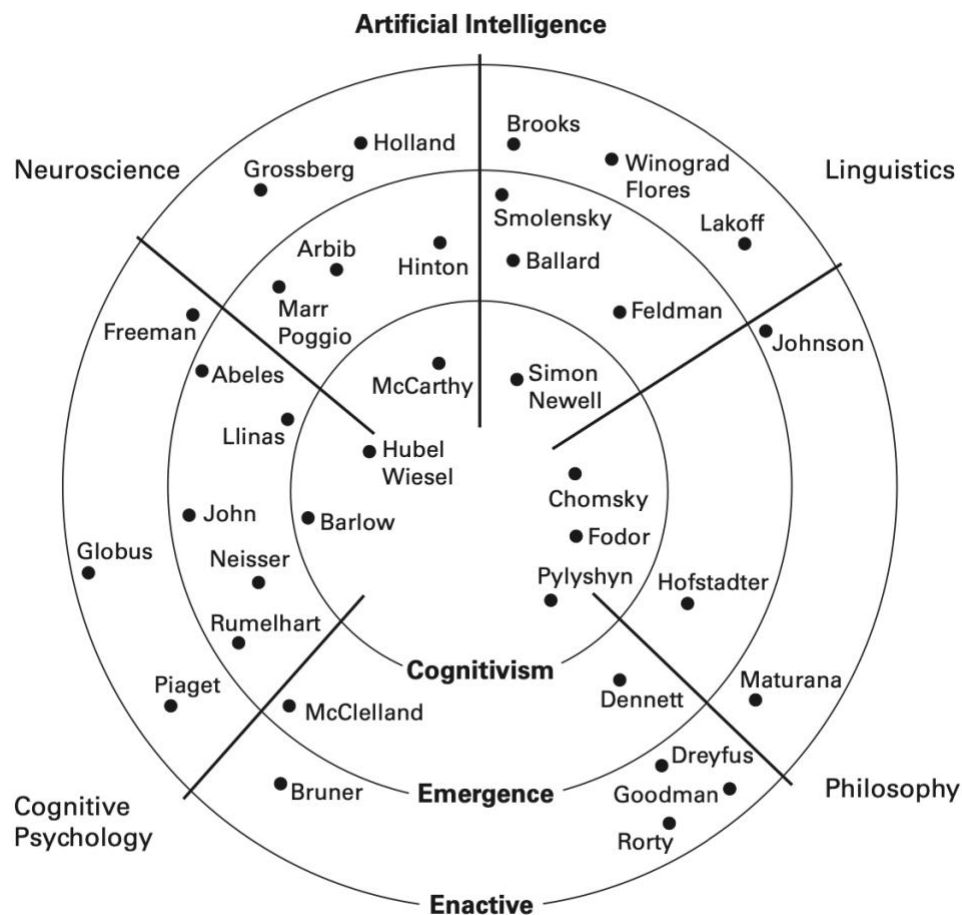


Figure 4: A polar map of the cognitive sciences, from Varela, Thompson and Rosch, *The embodied mind*, 7.

For the purposes of this thesis project, we can disregard emergence. This position, also known as connectionism, is related more to the field of artificial intelligence than to the philosophy of perception. Emergence is a position about *how* representations are instantiated. In this project, I am investigating *whether* perception involves

representations. Thus, the central disagreement arises between cognitivism and enactivism.

Looking at *Figure 4*, it is interesting to note the range of fields that are relevant to the representationalism question. Philosophers, neuroscientists, linguists and psychologists all investigate the question of whether cognition involves representation. Since this thesis project is empirically motivated, I will touch upon some of these fields, in addition to philosophy. It is also interesting to note the spectrum of positions between cognitivism and enactivism. Fodor and Pylyshyn are strongly cognitivist, but other philosophers take less polarised positions. In §2.3, I will discuss what it means to take a moderate position in the representationalism debate. There are, in fact, maximal and minimal definitions of ‘representation’.

## **2.3. What is representation?**

### **2.3.1. The function of carrying information**

Clark defines representations in terms of the function of carrying information.<sup>17</sup> However, representation involves something more than mere construal or mere correlation. Clark gives the example of the tides and the position of the moon; they are correlated, but neither *represents* the other.<sup>18</sup> So, crucially, a representation is an

---

<sup>17</sup> The notion of ‘function’, especially in the context of teleological explanations, is controversial. Clark largely avoids this controversy. His appeal to function describes representations as inner states which *operate so as to carry information and have operated as such* in the evolutionary history of the organism. The wider debate about teleological function is outside the scope of this thesis project. See: Allen and Neal, ‘Teleological notion in biology’.

<sup>18</sup> Clark, *Being there*, 146.

inner state (or process) which *has the function* of carrying specific types of information. This function is interpreted in an evolutionary, selective context. The inner state must have been selected, or must have been designed, or must have evolved for the purpose of carrying such information. For example, in the rat brain, the neurons in the posterior parietal cortex have the function of carrying information about the direction of the rat's head; thus, we can say that these neurons *represent* the rat's location in the environment (in a narrow way).<sup>19</sup>

Clark defines a representationalist 'story' in terms of a functional relationship between internal states and external conditions. A representationalist story involves internal states (or processes) which have the function of carrying specific types of information about external (or bodily) states of affairs.<sup>20</sup>

Pete Mandik gives a very similar definition of representation: 'neural events represent environmental events in virtue of having the function of carrying information about environmental events'.<sup>21</sup> Here, too, function is defined in terms of natural selection, so being naturally selected to carry information qualifies a neural state as being a genuine representation. For Mandik, this is the standard and accepted definition – 'a common way of thinking about representations, especially in neuroscientific contexts'.<sup>22</sup>

---

<sup>19</sup> Clark, *Being there*, 146.

<sup>20</sup> Clark, *Being there*, 147.

<sup>21</sup> Mandik, 'The neural accomplishment of objectivity', 456.

<sup>22</sup> Mandik, 'The neural accomplishment of objectivity', 456.

### 2.3.2. A continuum of representing

Once we define representation in terms of the function of carrying information, there appears to be a continuum of representational states. This continuum ranges from very simple inner states (about some feature of an outer event), to highly sophisticated inner states (that model abstract or imaginary entities). At the lower extreme, we find very simple cases of ‘adaptive hookup’. In these simple cases, the organism (or representational system) is hooked up to a narrow feature or property in the outer environment. Clark gives the example of a sunflower or a light-seeking robot moving towards the sun. Because the hookup is so primitive, ‘we gain little from treating the inner state as a representation’.<sup>23</sup> At the lower extreme of the continuum, these primitive cases barely qualify as representational states.

It is salient to consider these threshold cases. Clark is correct that what counts as representation is not something absolute; it doesn’t suddenly materialise, like a light being switched on. Rather, it shades in gradually from more primitive cases of correlation and information-bearing states, which – at some point – can be coherently called *representations*. On this continuum, Clark argues that ‘adaptive hookup’ gradually becomes ‘genuine internal representation’, once the internal state gains enough complexity and systematicity.<sup>24</sup>

At the upper extreme of the continuum, there are states which can carry information in the total absence of their objects, in the total absence of ‘target’ features. Humans and a handful of other species can engage in ‘complex imaginings’ and

---

<sup>23</sup> Clark, *Being there*, 147.

<sup>24</sup> Clark, *Being there*, 147.

counterfactual reasoning.<sup>25</sup> We can reason about ‘the distant’, ‘the non-existent’ and ‘the highly abstract’.<sup>26</sup> These problems necessitate the use of representational ‘stand-ins’; they are ‘representation-hungry’, because they appear unsolvable without such stand-ins.<sup>27</sup>

We are now entering intentionality territory, since we are discussing how internal states can *stand-in* for external states of affairs. The question of *how* such stand-ins are possible is a deep problem for naturalistic philosophy. That is the puzzle of naturalising intentionality. Thankfully, perhaps, it is not the subject of this thesis project. We are merely concerned with the fact *that* internal states stand-in for external states of affairs, not with *how* such stand-ins are possible.

Clark’s notion of a continuum has a clarifying effect. We can acknowledge that representations come in varying degrees of strength. Not all representations *stand-in* to the same degree; some are primitive hookups, while others are full-blown *reimaginings* of things that are not physically present. There are degrees of representing and degrees of mediation. At the start of §2, I introduced the representationalism debate using a crudely binary question: *is perception mediated or unmediated?* At the opposing extremes, cognitivists and enactivists do give binary answers to this question. However, there are also moderate positions, which say that perception is mediated *to some degree* – that perception involves *minimal* representations. That is, in fact, Clark’s position, which I will introduce in §3.

---

<sup>25</sup> Clark, *Being there*, 147.

<sup>26</sup> Clark, *Being there*, 175.

<sup>27</sup> Clark, *Being there*, 147.

### 3. Enactivism: Assessed

#### 3.1. Assessment overview

In §2, I introduced the representationalism debate: cognitivism, enactivism and moderate positions about representations. This taxonomy enables us to proceed to the critical discussion. In §3, I will analyse various enactivist accounts, all of which challenge representationalism to some degree. I will start with the most complete rejections of representations, then move to more moderate positions: (§3.2) Gibson’s *true anti-representationalism*; (§3.3) Varela, Thompson and Rosch’s *anti-foundationalist enactivism*; (§3.4) O’Regan and Noë’s *sensorimotor enactivism*; (§3.5) Clark’s *minimal representing*.

<i>Enactivism: Assessed</i>			
§	Heading	Author(s)	Anti-representationalist?
3.2	True anti-representationalism	Gibson	● ● ● ● ●
3.3	Anti-foundationalist enactivism	Varela, Thompson and Rosch	● ● ● ● ●
3.4	Sensorimotor enactivism	O’Regan and Noë	● ● ● ● ○
3.5	Minimal representing	Clark	● ● ● ○ ○

Table 2: Overview of §3 – Enactivism: Assessed

Individually, and together, these approaches challenge representationalism in various ways. These approaches all fit under the general umbrella of enactivism, either as a historical influence (Gibson) or as a contemporary offshoot of the movement.

Assessing these branches of enactivism will allow us to answer to **Sub-question A: *To what extent has enactivism challenged or influenced representationalism?***

In §3 and §4, I will argue in favour of **Conclusion A1: Perception is guided by both action and minimal representations**. I will provide philosophical and empirical evidence in favour of both limbs: the role of action and the role of minimal representations. This involves both a positive and negative assessment of the influence of enactivism. On the positive side, enactivists provide salient evidence that action guides perception. This supports their wider view that cognition is embodied and embedded. That is what we stand to learn about perception from studying enactivism. However, on the negative side, enactivists go too far in rejecting representations. I will argue that perception requires minimal representations. Moreover, I will argue that enactivism itself (namely, O'Regan and Noë's sensorimotor enactivism) *entails* minimal representations; otherwise, it commits an internal inconsistency. This leads to **Conclusion A2: Enactivism must be rehabilitated to include minimal representations**. I argue that enactivism can provide a robust account of perception, provided it is updated and 'rehabilitated' in this way.

### **3.2. Gibson: True anti-representationalism**

James J. Gibson's oeuvre begins in the middle of the twentieth century. His approach to vision, 'ecological psychology', rejects cognitivism and representationalism, and proposes a radical new theory of 'direct visual perception'. Gibson's core claim is that perception is completely unmediated – the direct pickup of information from the

environment. This founding manifesto is not fully accepted by the later enactivist philosophers, but certain elements of Gibson's theory are central to enactivism. Gibson emphasises bodily presence in the world, which is echoed by later theories about embodied and embedded cognition. There is a reciprocity between awareness of the body and awareness of the world: 'even an experience of pain is never wholly without some experience of the environment'.<sup>28</sup> In this way, the body and the world are not separable; they are dynamically and causally connected.

Gibson's theory is perhaps best illustrated by one 'false analogy to vision':

Now it is perfectly true that when an observer looks at a painting, photograph, sculpture, or model, he gets an *indirect* visual perception, a *mediated* experience, an awareness at *second hand*, of whatever is represented... Thus, there can be a direct perception of a man's portrait accompanied by an indirect perception of the man himself. The fallacy of the standard theories of perception consists of taking as a model for vision the kind of indirect visual perception that uses pictures as substitutes for things... We do not look at our retinal images and perceive the world in the way that we look at a portrait and perceive the sitter.<sup>29</sup>

For Gibson, visual experience is not mediated in any way. We see the world, not an image or representation of the world. We see the sitter, not the portrait. We have a direct relationship with perceptual objects in the world and we directly pick up information about those objects, without any epistemic mediator. Vision is not a

---

<sup>28</sup> Gibson, 'A theory of direct visual perception', 78.

<sup>29</sup> Gibson, 'A theory of direct visual perception', 88–9.



‘photographic process’, and the analogy between eye and camera is misleading. Instead, according to the theory of direct perception, vision is a process of ‘*sampling of the ambient array by the ocular system*’. Vision is ‘a process of exploration in time’.<sup>30</sup> This notion of ‘exploration’ is central to O’Regan and Noë’s later enactivism.

Fodor and Pylyshyn outline a neat criticism of Gibson’s theory. Gibson argues that certain properties of the ambient light array are directly picked up during perception. And yet, how does this information lead to perceptual knowledge? There must be some inferential mediation involved, which translates or otherwise interprets the raw information-in-the-light, producing knowledge. Gibson’s theory does not work, since the information in the ambient array *alone* is not enough; there must be some epistemic mediation.<sup>31</sup>

Fodor and Pylyshyn’s objection is astute. It may be that light, or even information-in-the-light, can be directly picked up, but how could *perceptual knowledge* or *perceptual experience* be directly picked up? Surely, knowledge is not ‘out there’ in the world, waiting for collection – like a soft drink in a vending machine. It seems reasonable to suppose that some, limited act of mediation (interpretation or inference) is required. Gibson’s realism is such that he supposes that the information ‘out there’ in the world is sufficient for perception. He is objecting to the idea of internal representations, and yet goes too far in rejecting all internal perceptual processes. Gibson’s theory of direct perception is unsuccessful, since it fails to account for the internal processes – on the part of the perceiver – that *interpret* information ‘out there’.

---

<sup>30</sup> Gibson, ‘A theory of direct visual perception’, 83–4.

<sup>31</sup> Fodor and Pylyshyn, ‘How direct is visual perception?’, 171.

Or, at least, it fails to show how objective information *alone* is sufficient for perceptual knowledge and perceptual experience. As we will see, very few other philosophers take such a radical line.

### 3.3. Varela, Thompson and Rosch: Anti-foundationalist enactivism

*The embodied mind*, by Varela, Thompson and Rosch, is a highly influential work in phenomenology and cognitive science. The authors make a ground-breaking case for ‘the enactive approach’, drawing on Buddhism and ‘enactive cognitive science’.<sup>32</sup> One central theme or thread in *The embodied mind* is the fundamental circularity between mind and world. Varela, Thompson and Rosch pose a chicken-and-egg question: ‘Which came first, the world or the image?’<sup>33</sup> They reject both the *chicken position* (realism), which says the world came first, and the *egg position* (idealism), which says that the image came first. Instead, they want to chart a ‘middle way’ between the two extremes. Just as with chickens and eggs, it makes no sense to think of mind and world as separate, with one more fundamental than the other. They exist in a perpetual feedback loop. Mind and world specify each other.<sup>34</sup>

Varela, Thompson and Rosch make an argument from Buddhism, from the *Madhyamaka* (‘middle way’) tradition. Human beings are consumed by *foundationalism* – ‘the search for an ultimate ground’.<sup>35</sup> The realist sees the *world* as the ultimate ground; mind-independent objects come before cognition. The idealist sees the *mind*

---

<sup>32</sup> Varela, Thompson and Rosch, *The embodied mind*, 9.

<sup>33</sup> Varela, Thompson and Rosch, *The embodied mind*, 172.

<sup>34</sup> Varela, Thompson and Rosch, *The embodied mind*, 172.

<sup>35</sup> Varela, Thompson and Rosch, *The embodied mind*, 144.

as the ultimate ground; object-independent minds come first, and they do the work of constructing the world. Realism and idealism share a common error in seeking *any* ultimate ground. Neither the mind nor the world is more fundamental.

Varela, Thompson and Rosch encourage a kind of radical 'anti-foundationalism'. Anti-foundationalism about knowledge says: there is no epistemological ground. ('We are like sailors who on the open sea must reconstruct their ship but are never able to start afresh from the bottom'.<sup>36</sup>) Varela, Thompson and Rosch go further than rejecting mere epistemological foundations. Their account entails an 'onto-epistemological' anti-foundationalism, since they reject any ultimate ground in *metaphysics* as well as in epistemology. Sebastjan Vörös and Alexander Riegler explain this onto-epistemological package: '*being* (ontology) and *knowing* (epistemology) co-determine/co-specify each other'.<sup>37</sup> For Varela, Thompson and Rosch, there is no prior, pregiven world. Rather, the world is *enacted* through collaboration and interaction with the organism. They want us to relinquish the need for any foundation, and to embrace the circularity of existence.

All this talk of feedback loops is made more precise by the term 'operational closure'.<sup>38</sup> Operational closure means that the results of a system's processes *are* those same processes. These are processes which, by their very nature, 'turn back upon themselves'.<sup>39</sup> Such systems are not *heteronomous* ('defined by external mechanisms of control') but rather *autonomous* ('defined by internal mechanisms of self-

---

<sup>36</sup> Neurath, 'Anti-Spengler', 199.

<sup>37</sup> Vörös and Riegler, 'A plea for not watering down the unseemly', 3.

<sup>38</sup> Varela, Thompson and Rosch, *The embodied mind*, 139.

<sup>39</sup> Varela, Thompson and Rosch, *The embodied mind*, 139.

organisation').<sup>40</sup> It is useful to use the term 'loopy' dynamics, borrowed from Daniel D. Hutto and Erik Myin.<sup>41</sup> Hutto and Myin provide a witty (and presumably original) formulation for the dynamics of mentality. They argue that cognition arises from dynamic interactions which are 'loopy', as opposed to linear.<sup>42</sup> According to loopy dynamics, there is no clear inner/outer distinction or input/output distinction. Instead, cognition arises from recursive feedback loops, coming from all parts of the brain, body and environment.

*Loopiness* is the distinguishing feature of Varela, Thompson and Rosch's enactivism. The system does not simply construct a world in the form of an internal model (idealism). Rather, the system enacts the very world itself. Here, the word 'enacts' (hence, 'enactivism') refers to this dynamical interaction between perceiver and environment that gives rise to perceptual objects. Perceptual objects are enacted by organisms in specific contexts; they do not exist prior to, or independently of, cognition, and yet they are not mere constructs of the mind. Instead, they are 'enacted' through physical interaction with the environment. In this way, the system and the world itself mutually specify each other, in a loopy way. The dynamics of the system and the dynamics of the world are so wrapped up in one another that they are inseparable. Outside the loop, there is no mind and no world of which to speak. In other words, there is no mind-independent world and there are no world-independent minds.

---

<sup>40</sup> Varela, Thompson and Rosch, *The embodied mind*, 140.

<sup>41</sup> Hutto and Myin, *Radicalizing enactivism*, 6.

<sup>42</sup> Hutto and Myin, *Radicalizing enactivism*, 6.

These are weighty, metaphysical claims. The enactivism espoused by Varela, Thompson and Rosch is a far-reaching proposal about the metaphysical arrangement between perceiver and world. Such claims – about the metaphysical nature of mentality and the metaphysical nature of the world – are of central importance to this thesis project. What should we make of Varela, Thompson and Rosch’s metaphysical proposal? I must leave this cliffhanger for the later sections on metaphysics and realism. I will analyse Varela, Thompson and Rosch’s onto-epistemological anti-foundationalism – along with its commitment to anti-realism – in §7.3.3. For the time being, I will press ahead with my analysis of the representationalism debate, and the various enactivist approaches to representation in perception.

### **3.4. O’Regan and Noë: Sensorimotor enactivism**

#### **3.4.1. Sensorimotor contingencies**

In 2001, J. Kevin O’Regan and Alva Noë published, ‘A sensorimotor account of vision and visual consciousness’, which has been described as a ‘seminal publication’.<sup>43</sup> In this paper, O’Regan and Noë set out their own brand of sensorimotor enactivism. Their ‘sensorimotor contingency theory’ is a new and innovative approach to vision. They reject the orthodox view that vision requires detailed internal representations of the outside world. Instead, they conceive of vision as ‘an *exploratory activity*’. In

---

<sup>43</sup> Bishop and Martin, ‘Contemporary sensorimotor theory: A brief introduction’, 1.

O'Regan and Noë's own words, '*vision is a mode of exploration of the world that is mediated by knowledge of what we call sensorimotor contingencies*'.<sup>44</sup>

Sensorimotor enactivism involves a change of focus – *away* from retinal images and internal models, and *towards* the agent's actions, which cause law-like changes in the agent's perceptions.<sup>45</sup> Perceptions are formed through actions. Specifically, an agent's perceptual experience arises when that agent interacts with the world; this is the process which *enacts* vision. When an organism moves through the world, it moves its body through a physical landscape, engaging directly with the physical objects of perception. The organism *explores* this landscape – moving nearer to, or further from, or around, physical objects.

The notion of 'sensorimotor contingencies' is central to O'Regan and Noë's theory. Sensorimotor contingencies are defined as '*the structure of the rules governing the sensory changes produced by various motor actions*'.<sup>46</sup> There are indeed 'rules' that make vision different from audition or touch (or any other sensory modality). Each modality is governed by its own rules. O'Regan and Noë draw attention to these lawlike contingencies, arising from interactions between physics and physiology. Such rules can be abstracted and induced, and this is precisely what vision entails; seeing involves understanding and mastering the laws of sensorimotor contingency.

O'Regan and Noë use the example of a missile guidance system. The system is 'tuned to' the relevant laws. Moreover, it has 'mastery over' these laws – accelerating or dipping or banking in order to follow the target airplane. There is a masterful

---

<sup>44</sup> O'Regan and Noë, 'A sensorimotor account of vision and visual consciousness', 940.

<sup>45</sup> Bishop and Martin, 'Contemporary sensorimotor theory: A brief introduction', 1.

<sup>46</sup> O'Regan and Noë, 'A sensorimotor account of vision and visual consciousness', 941.

coupling of input and output; the missile redirects its body in response to the moving image of the airplane ahead of it.<sup>47</sup>

Let's return to the example of Tiger Woods, which I discussed in §2.2.3. The most important feature of O'Regan and Noë's theory is the *contingent* relationship between perception and action (hence, 'sensorimotor contingencies'). Tiger knows: *if* I do this, *then* I will perceive that. *If* I swing the driver at this precise angle, *then* I will perceive the golf ball whizzing down the fairway. Perception requires 'mastery' of these sensorimotor contingencies, in masters such as Tiger Woods. For O'Regan and Noë, perception is grounded in action and sensorimotor contingencies, rather than representations. However, there is already one issue that is crying out. When Tiger knows about the relevant contingencies (if I swing the club, then I will see the ball moving), must he not *represent* these contingencies – or, at least, represent his expectations about contingencies? O'Regan and Noë claim to be anti-representationalist, but perhaps their account entails representations. I will analyse this issue, at length, from §3.4.3 to §3.4.6.

In recent years, Noë has preferred the term 'actionism' rather than 'enactivism'. Noë flags this change of terminology in a 2015 paper, 'Concept pluralism, direct perception, and the fragility of presence': 'The present study takes its starting point from the enactive or sensorimotor, or, as I now prefer to call it, the actionist approach to perception and perceptual consciousness'.<sup>48</sup> It is useful to note this refinement in terminology by Noë. However, since this thesis project is focused more widely on the

---

<sup>47</sup> O'Regan and Noë, 'A sensorimotor account of vision and visual consciousness', 943.

<sup>48</sup> Noë, 'Concept pluralism, direct perception, and the fragility of presence', 1.

enactivist tradition, and since a great deal of secondary literature still uses the enactivist label for Noë's work (as Noë did himself until recently), I will continue to use the term enactivism, rather than actionism.

### **3.4.2. Action, exploration and the world as an outside memory**

There is a special role for action in O'Regan and Noë's theory. Action grounds perception, and action is also the way that an organism *unlocks* information from its surroundings. Information in the environment can be accessed through various actions:

Perceivers... squint, lean forward, reach for their glasses, tilt their reading material to the light, and casually walk to the window to get glimpses. They act as though they believe that the detail is there, in front of them, and that to acquire detail, to bring it into consciousness, they need to act.<sup>49</sup>

Here, Noë and O'Regan depict their claim that perception involves active exploration of the environment. Vision is made possible by these bodily activities – squinting, reaching, tilting, walking. We might recall the bodily activities of Tiger Woods on the first tee. He must scan the fairway, shading his face from the sun, and then execute a precise stroke using a carefully selected golf club. These are the actions that Woods takes to enact the golf course in his visual experience.

---

<sup>49</sup> Noë and O'Regan, 'On the brain-basis of visual consciousness', 577.



This special role for action – enabling access to information in the environment – is summed up by O'Regan's theory of 'the world as an outside memory'. It was first outlined in his paper, 'Solving the "real" mysteries of visual perception', and later echoed in various other texts by O'Regan and Noë. O'Regan begins by listing apparent 'defects' in the human eye: inverted structure; 'blind spot'; nonuniformity of cones in the retina; nonuniformity of rods; smearing and displacement caused by eye movements.<sup>50</sup> He homes in on two of the "'real" mysteries': how we are unaware of the blind spot, and how we are unaware of the saccadic perturbations. He rejects the orthodox explanations involving 'compensatory mechanisms', which presuppose some kind of 'internal screen'. Instead, O'Regan puts forward his own, radically different explanation. The outside world serves as an 'external memory' which can be accessed at will via eye movements. Thus, there is no need for detailed internal representations; the world serves as its own best model.<sup>51</sup>

Noë illustrates the world as an outside memory theory in terms of an Internet metaphor. Just like an Internet user, an organism off-loads its processing demands by storing information remotely, rather than locally. When an Internet user loads a website, all the detail is present (or appears as if present) on the user's desktop machine. However, this information is not present *locally* on the user's machine. Instead, it is stored *remotely*, via online servers, and it is accessed virtually. Noë rejects the cognitivist idea that information about the environment is reduplicated onto the local hard drive: 'there's no need to *re-present* the world on one's own internal

---

<sup>50</sup> O'Regan, 'Solving the "real" mysteries of visual perception', 462.

<sup>51</sup> O'Regan, 'Solving the "real" mysteries of visual perception', 463.

memory drive’.<sup>52</sup> Instead, the world is right there, containing all the relevant information about the world as if on a remote server. Since the organism is ‘networked’ to the world, the information can be directly accessed without *re*-presentation.<sup>53</sup>

Noë’s Internet metaphor is a neat illustration of how sensorimotor enactivists conceive of representations; they are rejected as unnecessary *reduplication* of information that is already present in the world. The most interesting feature of this claim is a tacit, metaphysical assumption. O’Regan and Noë appear to make a naively realist presupposition – that the world is straightforwardly *there*, to be accessed at will. This tacit assumption is by no means obvious, and it will be a matter for investigation in §6 and §7.

### 3.4.3. Internal processing and sensorimotor knowledge

In the following subsections, I will lay out the case for **Conclusion A2: Enactivism must be rehabilitated to include minimal representations**. I will demonstrate that O’Regan and Noë’s enactivism entails minimal representations, despite their explicit rejection of representations. This is first evidenced by O’Regan’s criticism of Gibson’s direct perception.

O’Regan discusses the direct perception versus indirect perception debate. At first glance, his theory appears to endorse ‘direct pickup’ of environmental stimuli, *qua* Gibson, since the world is an external memory which can be accessed at will.

---

<sup>52</sup> Noë, *Action in perception*, 50.

<sup>53</sup> Noë, *Action in perception*, 50.

However, O'Regan rejects Gibson's radical claims about direct perception, writing that his external memory theory should be 'distinguished from a radical Gibsonian viewpoint'.<sup>54</sup> Under Gibson's theory, internal representations play *no* role, whereas, under O'Regan's theory, there is *some* 'processing' of information. Granted, the world is an external memory store for information, and that information can be directly accessed, but it is still *processed* in some way. The stimuli that we see are the ones that are processed or integrated via 'the appropriate cognitive operations'.<sup>55</sup>

O'Regan's position here is slightly puzzling. In his introduction, he rejects the notion of an 'epistemic mediator' (any internal representation) since the world is straightforwardly accessible. Yet, later on, he accepts that there are *some* cognitive operations involved in processing stimuli. He is vague as to which cognitive operations, but we can speculate that it involves *more* than Gibsonian direct perception, and *less* than a full-blown internal image. So, O'Regan seems to favour a type of minimal representationalism, but remains vague as to the details. Noë and O'Regan (separately and together) write forcefully about rejecting all internal representations, and yet their framework is not completely anti-representationalist. They accept some internal processing, some minimal representing – in contradiction to their sweeping claims about being anti-representationalist. This is an intriguing contradiction.

What type of minimal representations are entailed by O'Regan and Noë's enactivism? We get some indication from O'Regan's later work, and from Noë's

---

<sup>54</sup> O'Regan, 'Solving the "real" mysteries of visual perception', 473.

<sup>55</sup> O'Regan, 'Solving the "real" mysteries of visual perception', 473.

monograph, *Action in perception*. Noë writes: ‘whether you call it inference or not, perception is mediated by a complex cognitive process whereby we recover facts about the layout’.<sup>56</sup> The perceiver must understand and reconstruct facts about the environment. The perceiver must represent their contingent relationship to the environment. O’Regan and Noë’s enactivism requires that the perceiver has *sensorimotor knowledge* of these contingencies:

I have been arguing that vision depends on sensorimotor knowledge... It is this knowledge of the way sensory stimulation varies as a function of movement that is the basis of our ability to have world-presenting sensory experience.<sup>57</sup>

Noë writes that sensorimotor knowledge is ‘practical’, *not* propositional. It is the knowledge that accompanies mastery of the laws of sensorimotor contingencies: knowledge of the if/then relationships connecting movement and vision.<sup>58</sup> The perceiver must know how movements will affect the appearance of stimuli in the visual field. Crucially, this conception of sensorimotor knowledge entails some form of representation. The perceiver (like Tiger Woods) must have a thorough knowledge of the relevant contingencies: if I do this, then I will perceive that. Tiger must *know* and *represent* these contingencies. For O’Regan and Noë’s theory to work, Tiger (or any perceiver) must have some representation about what he will perceive, if he acts.

---

<sup>56</sup> Noë, *Action in perception*, 104.

<sup>57</sup> Noë, *Action in perception*, 117.

<sup>58</sup> Noë, *Action in perception*, 117.

If perception is indeed guided by sensorimotor knowledge, then representations cannot be fully eliminated. At least some minimal representing is required. What form do these representations take? It appears that there are two candidates: (§3.4.4) *representing expectations* and (§3.4.5) *representing an index*.

### 3.4.4. Representing expectations

What is sensorimotor knowledge, and where does it come from? It is reasonable to suppose that this knowledge is grounded in the agent's *expectations* and *anticipations* about their own actions.<sup>59</sup> Noë seems to endorse this minimal form of representing – or, at least, his account would not be possible without this minimal form of representing. Noë writes that ‘the work of the enactive approach is done by perceivers’ *expectations* of the sensory effects of movement’.<sup>60</sup> In his view, agents achieve a type of practical, sensorimotor knowledge and it makes sense to interpret this knowledge as the anticipations and expectations of the agent.

Sensorimotor theory appears to rely on a type of visual sampling. The world is used as its own best model – removing, ‘at a stroke’, the need to construct or store or update any internal models of the world.<sup>61</sup> Instead, the organism has a direct encounter with the world; it engages in ‘active exploration of the environment’ through ‘a series of [saccadic] movements’, which serve to confirm (or correct) previous predictions.<sup>62</sup> Indeed, O’Regan and Noë talk about sampling: ‘visual

---

<sup>59</sup> Hutto and Myin, *Radicalizing enactivism*, 26–27.

<sup>60</sup> Noë, *Action in perception*, 119.

<sup>61</sup> Bishop and Martin, ‘Contemporary sensorimotor theory: A brief introduction’, 2.

<sup>62</sup> Bishop and Martin, ‘Contemporary sensorimotor theory: A brief introduction’, 2.

exploration provides ways of sampling' the three-dimensional world.<sup>63</sup> The organism methodically (saccadically) scans the environment, checking whether the *actual* stimulus is the same as the *anticipated* stimulus. It is a trial and error process: scanning, matching, adjusting.

Is there empirical support for the idea that perception involves active sampling? It is instructive to turn to the relevant neuroscience research. There is experimental evidence which aligns with the (broad) description of such a saccadic sampling process. Didier Flament and his team conducted fMRI research to study visuomotor learning.<sup>64</sup> In this experiment, subjects used a joystick to superimpose a cursor onto visual targets. For one group, the joystick/cursor relationship was inverted and they proceeded to learn the new task, while the research team captured brain images.

The team found changes in cerebellar activation, and these activation levels were highest during the early stages of learning the new task. Moreover, as soon as performance improved, there was a decrease in cerebellar activation. Indeed, there was a parallel (inverse) relationship between performance and activation.<sup>65</sup> Flament and his co-authors conclude that inaccurate motor performance causes the cerebellum to activate more strongly. This supports the hypothesis that the cerebellum is involved in detecting and correcting visuomotor errors.<sup>66</sup>

---

<sup>63</sup> O'Regan and Noë, 'A sensorimotor account of vision and visual consciousness', 941.

<sup>64</sup> Flament et al., 'Functional magnetic resonance imaging of cerebellar activation during the learning of a visuomotor dissociation task.

<sup>65</sup> *Ibid.*, 210–221.

<sup>66</sup> *Ibid.*, 224.

It might be possible to extrapolate from this study, tracing the implications for enactivism. There is indeed evidence that the visuomotor system *learns*, through a process of sampling and error correction. Activation in the cerebellum indicates an error signal, enabling the organism to adjust – improving visuomotor performance. This research certainly explains a puzzle surrounding conscious versus unconscious experience. We consciously experience events which are unpredictable – triggering an error signal. Meanwhile, routine events are perceived unconsciously, since the *actual* stimulus aligns with the *anticipated* stimulus.

What does this mean for the enactivist's critical claim, that a model is redundant? On the contrary, the existence of a 'prediction' requires some representation or model of *what is anticipated*. Naturally, the prediction error would be generated when the outcome deviates from that model; a system of trial and error learning requires such predictions and corrections. So, it appears that Noë endorses a type of minimal representationalism – comprised of the agent's expectations. This creates an inconsistency in Noë's treatment of representations. As the French would put it, *on ne peut pas avoir le beurre et l'argent du beurre* (you can't have the butter and the money for the butter).<sup>67</sup> Having the butter and having the money for the butter are mutually exclusive. Equally, rejecting representations and appealing to representations are mutually exclusive. Noë's inconsistency about representations amounts to wanting both the butter and the money for the butter. Thus, enactivism is in need of rehabilitation.

---

<sup>67</sup> The French equivalent of 'you can't have your cake and eat it too', arguably a sharper and more astute formulation.

### 3.4.5. Representing an index

In addition to representing expectations, O'Regan and Noë's theory appears to require a second type of representation. This second type of representation relates to an egocentric 'I am here' marker on the 'external memory' model of the world. O'Regan's notion of the world as an external memory is a core part of O'Regan and Noë's rejection of representations, and yet an external memory is effectively useless without some internal index. The subject must represent some index of what is contained in the memory drive, or else they would have to resort to 'random search'. The mere idea that the world is continuously available 'provides no comfort to someone who has lost her keys'.<sup>68</sup>

Let's imagine that the perceiver, Bethany, has lost her keys. According to O'Regan's radical claim, Bethany has no internal model of the world. The world is continuously available, so she has no need to model and re-present the world. However, in reality, this would leave Bethany unable to engage with her environment. She needs an index of how the world relates to her. For one, she needs an egocentric marker, to locate herself on the 'external' model of the world. She also needs an index of the objects and entities that are relevant to her: food, predators, offspring... socks, earphones, car keys. Without this index – without some guidance on *where* to look – the world is uniformly featureless. Every object is uniformly relevant, or uniformly irrelevant, providing no information whatsoever. Therefore, Bethany needs to

---

<sup>68</sup> Wallis and Wright, 'Enactivism's vision', 259.



represent an egocentric index of which entities are relevant, relative to where she is. I will return to this notion of an egocentric index in §5.2.4.

The need for this index points to a wider issue of representation and resources. Enactivists' boast that their framework has explanatory simplicity. However, much to the contrary, it appears that perception requires *large* amounts of 'stored know-how in the brain'.<sup>69</sup> O'Regan's account of object recognition 'seems to require that subjects know a great deal of SMCs [sensorimotor contingencies]'.<sup>70</sup> And, here, there is a contradiction. On the one hand, enactivists reject cognitivism as being too resource-heavy (the resources needed to construct detailed, internal models). On the other hand, they presuppose significant stores of SMC-related knowledge. This is the same inconsistency, discussed in relation to representing expectations. O'Regan and Noë claim to be anti-representationalist, and yet their theory requires representational resources. Enactivists argue that perception is direct, yet the perceiver must possess a great deal of *prior* sensorimotor knowledge. Certainly, enactivism is in need of some rehabilitation.

### 3.4.6. Rehabilitating enactivism

Enactivists do not take a consistent stance on representationalism. They sometimes come across as radically anti-representationalist, making 'headline grabbing' statements: vision is a process of exploration, not a process of internal *re-*

---

<sup>69</sup> Wallis and Wright, 'Enactivism's vision', 260.

<sup>70</sup> Wallis and Wright, 'Enactivism's vision', 261.

presentation.<sup>71</sup> Then, they appear to backtrack, with more moderate and sober statements: ‘we argue *not* that there are no representations, but that the category “representation” should be demoted within the context of visual theory’.<sup>72</sup> One of the most overtly *pro*-representationalist statements reads: ‘The work of the enactive approach is done by perceivers’ *expectations* of the sensory effects of movement’.<sup>73</sup>

That is the inconsistency in the enactivist framework. They set out to challenge representationalism. Yet, at the same time, their theory *entails* at least two forms of representing. First, their theory is only coherent if subjects are able to represent their expectations or anticipations arising from a certain action; the outcome then confirms or disabuses the prediction. Second, the perceiver must represent an index of the relevant features in the environment, relative to the organism’s physical location and physical body. It appears that O’Regan and Noë fail to anticipate this objection, since they don’t address it anywhere in their oeuvre. And this inconsistency must be addressed, if enactivism is to be a robust account of perception.

In §3, we have already seen mounting evidence in favour of **Conclusion A1: Perception is guided by both action and minimal representations**. On the one hand, O’Regan and Noë’s description of perception as guided by sensorimotor knowledge is astute. Moreover, it aligns with the neuroscientific evidence reported by Flament and his team. They, too, describe visuomotor perception as guided by a trial-and-error process, where the error signal is informed by knowledge about what is anticipated. On the other hand, there is evidence that representations cannot be eliminated from

---

<sup>71</sup> Wallis and Wright, ‘Enactivism’s vision’, 261.

<sup>72</sup> Noë, ‘What does change blindness teach us about consciousness?’, 218.

<sup>73</sup> Noë, *Action in perception*, 119.

perception. In §3.5, I will show how an action-centred view of perception is compatible with action-oriented representations. This will lead to **Conclusion A2: Enactivism must be rehabilitated to include minimal representations**. It is still possible for O'Regan and Noë to reject the other extreme of representationalism, with rich, action-neutral models of the world. Yet they must embrace minimal representing as a core part of their programme. After this gentle rehabilitation, sensorimotor enactivism becomes a cogent and promising account of perception.

### 3.5. Clark: Minimal representing

#### 3.5.1. An ecumenical approach to representationalism

Throughout §3.4, I argued that O'Regan and Noë's account amounts to a kind of minimal representationalism, despite their claims to be anti-representationalist. Moreover, once we realise that their framework entails minimal representing, it becomes a cogent account of perception. I will conclude §3 – my assessment of enactivism and representationalism – by analysing Andy Clark's book *Being there: Putting brain, body and world together again*. Clark is explicitly in favour of minimal representationalism: an ecumenical approach to representationalism. This is just the sort of position that aligns with our rehabilitated enactivism. By bringing Clark on board, we can build an even stronger case for perception as guided by minimal representing.

Clark takes a moderate stance in the representationalism debate. He takes a positive view of contemporary neuroscience, with its *combination* of radical and

traditional elements. He writes: 'it is representationalism and computationalism stripped of all excess baggage, and streamlined'.<sup>74</sup> This is indeed Clark's vision for the field: *Representationalism, streamlined*. He wants to save the notions of computation and representation, updating them in light of recent advances. He believes in the salience of some of these radical elements, but also recognises the robustness of the traditional approach to cognition.

Clark is embarking on a project of 'integration and reconciliation'.<sup>75</sup> A successful account of cognition must achieve a careful balancing act between old and new. Clark advocates an 'ecumenical position'; minds may be *both* 'essentially embodied and embedded', *and* – at the same time – '*still* depend crucially on brains which compute and represent'.<sup>76</sup> There is no inherent contradiction between the embodied and embedded approach, on the one hand, and computation and representation, on the other. That is the harmony that Clark preaches. That is his ecumenical sermon. I will argue that Clark's account is a perfect candidate for a rehabilitated enactivism, one that embraces both minimal representing and the role of action.

### 3.5.2. Minimal versus maximal representations

Clark's middle ground is a kind of 'minimal representationalism'.<sup>77</sup> The debate about representationalism has been overblown by a misconception about the *minimal* and

---

<sup>74</sup> Clark, *Being there*, 142.

<sup>75</sup> Clark, *Being there*, 142.

<sup>76</sup> Clark, *Being there*, 143.

<sup>77</sup> Clark, *Being there*, 174.

the *maximal* standpoints. Some view representations as maximal (detail, action-neutral models) as opposed to minimal (partial, personalised, action-oriented models).<sup>78</sup> The radicals tend to reject the maximal view without considering the minimal view, which is a perfectly coherent alternative. There might even be a trace of a straw man argument contained in the rejection of the maximal view only; radicals reject this opposing extreme (full-blown, action-neutral models), without considering the coherent middle ground (partial, action-oriented models). On Clark's assessment, minimal representationalism is the correct approach to understanding cognition.

Clark argues that Gibson has conflated the minimal and maximal views. Indeed, Gibson objects to the maximal view of representations, as action-neutral models. Gibson does not consider the minimal view, invoking action-oriented models. *Only* the maximal view is in conflict with the embodied and embedded approach to cognition. It is perfectly consistent to accept embodied and embedded approach *alongside* the minimal view of representations. Thus, Gibson is addressing a non-issue.<sup>79</sup>

Again, there is a faint shadow of a straw man in the radicals' rejection of representations. Clark is correct that Gibson (and other radicals) reject the *maximal* notion of full-blown, action-neutral models. They seem to ignore the more reasonable middle ground. It does appear possible to account for perception as grounded in action and as guided by action-oriented models. To that extent, Clark's ecumenical sermon is highly persuasive.

---

<sup>78</sup> Clark, *Being there*, 174.

<sup>79</sup> Clark, *Being there*, 172.

Clark's position (explicit minimal representationalism) pairs usefully with O'Regan and Noë's framework (implicit minimal representationalism). As I argued in §3.4, there is strong evidence – both conceptual and empirical – that perception requires minimal representations. Even if perception is guided by action, perceivers must still represent their expectations arising from potential actions. Perceivers must also represent an index of the relevant environment features. This aligns with Clark's notion of partial, personal, action-oriented models.<sup>80</sup> Now, we are building up an account of perception that accords with the evidence. Thus, I argue for **Conclusion A2: Enactivism must be rehabilitated to include minimal representations.** O'Regan and Noë's account must be updated to include Clark's notion of action-oriented models.

In §4, I will consider some empirical evidence for and against enactivism. This will shed more light on the role of action and the role of representation in perception. I will also provide more evidence and arguments for **Conclusion A1: Perception is guided by both action and minimal representations.**

---

<sup>80</sup> Clark, *Being there*, 174.

## 4. Enactivism: Cognitive science and empirical evidence

### 4.1. Evidence for enactivism

#### 4.1.1. Perception does not involve detailed internal representations

O'Regan and Noë discuss several examples supporting the claim that *perception does not involve detailed internal representations*, including: change blindness; the blind spot; and why the world appears stable despite eye movements.<sup>81</sup> I will briefly discuss one of these examples: trans-saccadic fusion (visual stability). Vision scientists have long investigated the problem of why the visual world appears stable, despite our constant and rapid eye movements. A great deal of experiments assumed the existence of an internal screen, which means that successive snapshots must be matched to a location on the screen, guided by an 'extraretinal signal'.<sup>82</sup>

O'Regan and Noë challenge this whole research paradigm. If there is no internal screen, then there is no need to measure or verify an accurate extraretinal signal. Indeed, David E. Irwin provides experimental evidence that there is no 'trans-saccadic fusion', rebutting the notion of an internal screen.<sup>83</sup> Moreover, without the notion of an internal screen, the 'problem' of visual stability becomes a 'non-problem'. The subject does not have to construct a composite patchwork of individual snapshots. Instead, the world acts as an outside memory.<sup>84</sup> O'Regan and Noë suggest that visual stability ('stationarity') comes from the *knowledge* of how eye movements

---

<sup>81</sup> O'Regan and Noë, 'A sensorimotor account of vision and visual consciousness', 949.

<sup>82</sup> O'Regan and Noë, 'A sensorimotor account of vision and visual consciousness', 949.

<sup>83</sup> O'Regan and Noë, 'A sensorimotor account of vision and visual consciousness', 950.

<sup>84</sup> O'Regan and Noë, 'A sensorimotor account of vision and visual consciousness', 950.

and bodily movements relate to objects in the world. Thus, it is mastery of the rules of sensorimotor contingency which allows an organism to perceive a stable world.<sup>85</sup>

#### 4.1.2. Action guides perception

O'Regan and Noë discuss several examples supporting the claim that *action guides perception*, including: sensorimotor adaptation, sensory substitution and synaesthesia.<sup>86</sup> One of the most interesting sets of experiments involves the 'inversion of the visual world'.<sup>87</sup> Subjects are given optical apparatus which inverts the retinal image, either upside-down or left-right. After roughly two weeks, they start to perceive the world as 'normal' again, despite their image of the world now being inverted. O'Regan and Noë discuss some intriguing effects that are reported during this two-week period of adaptation.

During adaptation, subjects experience a 'fragmentation' of perception, and a 'dependence on context and task'. Flames must go up, so a candle flips when it is lit, whereas coffee must pour downwards, so a mug flips when coffee is being poured. Another study reports that two adjacent heads, one upright, the other inverted, were *both* perceived as upright.<sup>88</sup>

O'Regan and Noë discuss these implications. Cognitivism cannot easily account for this fragmentation of perception. If perception is guided by internal representations, why would one object be orientated correctly, while another (*'sharing*

---

<sup>85</sup> O'Regan and Noë, 'A sensorimotor account of vision and visual consciousness', 949.

<sup>86</sup> O'Regan and Noë, 'A sensorimotor account of vision and visual consciousness', 949.

<sup>87</sup> O'Regan and Noë, 'A sensorimotor account of vision and visual consciousness', 953.

<sup>88</sup> O'Regan and Noë, 'A sensorimotor account of vision and visual consciousness', 953.



*the same retinal location'*) be orientated incorrectly? Under cognitivism, there is no mechanism which would explain or justify this fragmentation. However, under the enactivist view, action guides perception, without requiring internal models. This provides a ready explanation for how perception could be fragmented; vision is guided by action, so it can be segmented by different actions.<sup>89</sup> Different activities (reading versus depth perception versus facial recognition) require different faculties, and this could compartmentalise different stimuli. So, the strange findings from the inverted world experiments appear to support an action-centric account of perception.

O'Regan and Noë also investigate evidence that one sense modality can be 'rerouted' to another. For instance, for subjects who can read Braille, 'tactile processing is "rerouted" to occipital visual cortex'.<sup>90</sup> This supports the enactivist view that the brain supports motor functioning, which in turn gives rise to *perception* in the broad sense. The exploration of the environment depends on how the organism interacts with that environment, using their own particular sensory apparatus. While the brain supports these faculties, the process of 'seeing' involves so much more than mere brain activity.

---

<sup>89</sup> O'Regan and Noë, 'A sensorimotor account of vision and visual consciousness', 953.

<sup>90</sup> O'Regan and Noë, 'A sensorimotor account of vision and visual consciousness', 968.

## 4.2. Evidence against enactivism

### 4.2.1. The significance of sensorimotor contingencies

Some of the most widely discussed evidence in favour of enactivism involves visual adaptation or ‘prismatic adaptation’. These include the famous upside-down goggle experiments, where the subject’s retinal image is inverted. One of O’Regan and Noë’s central claims is that qualitative experiences arise due to sensorimotor contingencies (SMCs), so they can explain experiential blindness in terms of SMCs. O’Regan and Noë believe that inverted world experiments support their case, since enactivism can explain the experience of fragmentation.

Yet, there is a glaring problem with the enactivist case. The upside-down goggles invert the retinal image, completely disrupting and – literally – *upending* the relevant SMCs. *If* qualitative experience was so closely governed by SMCs, *then* surely such a disruption would make visual experience impossible. But this is not the case.<sup>91</sup> The subject is able to continue having visual experiences, and ultimately (after about two weeks) feels that the world looks normal again. This would seem to undermine, and even disprove, the claim that visual experience is closely governed by knowledge of SMCs.

Wallis and Wright conclude with several severe verdicts. They write that enactivists claims amount to: ‘deeply flawed interpretations’ of the literature; ‘ignorance of vast amounts’ of the literature; ‘failure to consider alternative

---

<sup>91</sup> Wallis and Wright, ‘Enactivism’s vision’, 263–4.

explanations'; and 'failure to correctly characterise predictions'.<sup>92</sup> These remarks are heavy hitting, with Wallis and Wright holding nothing back.

Let's consider how O'Regan and Noë would respond to this interpretation of the inverted world experiments. They would argue that the findings do not disprove the claim that visual experience is closely governed by knowledge of SMCs. After two weeks, the subject sees the world as 'normal', but in the *meantime*, there is a period of adaptation when perception is fragmented. This appears to add support for enactivism. The relevant SMCs have *indeed* been upended. The subject *does* experience a disruption to visual experience, until such time that they have re-learned the new SMCs of this new, inverted world. Thus, the evidence may still support O'Regan and Noë's action-centred account of vision.

### 4.3. Closing comments on representationalism

In §2, §3 and §4, I considered a range of philosophical and empirical evidence relating to enactivism. We are now in a position to formulate an answer to **Sub-question A: *To what extent has enactivism challenged or influenced representationalism?*** In a nutshell, enactivism poses a distinct and partially successful challenge to representationalism. In some ways it is astute, and in other ways it is flawed. It is certainly true that action guides perception, yet this does not entail the more radical enactive claims in favour of anti-representationalism. It is possible to have action-centred perception alongside representationalism. The evidence appears to support

---

<sup>92</sup> Wallis and Wright, 'Enactivism's vision', 266.

more moderate forms of representing – Clark’s minimal representationalism. Thus, the evidence leads us to **Conclusion A1: Perception is guided by both action and minimal representations.**

Remarkably, O’Regan and Noë’s theory appears to largely align with moderate representationalism. They claim to be radical enactivists, dispensing with representations, and yet their theory requires a form of minimal representing. This is not necessarily fatal, since it points to the salience of minimal representing across many frameworks. Representing minimal features of the world (expectations or an index) allows the organism to navigate the world successfully – still without detailed, action-neutral models. Thus, we reach **Conclusion A2: Enactivism must be rehabilitated to include minimal representations.**

This thesis project is not aimed at giving a definitive answer to the representationalism question. Instead, the purpose of this thesis is to unpack the enactivist’s claims about perception and, in particular, the metaphysical implications of enactivism. So, for our purposes, what remains to be shown is how an action-oriented model can be ‘personal’. I will take great care to investigate this issue in §5. I will argue that an organism’s ‘personal’ model of the environment should be understood in terms of *egocentricity*. An organism perceives the world in a way that is selective and goal-relevant, relative to their physical body and physical location in the environment. This account of egocentricity is the next major piece in my overall account of perception, and it will also shed more light on the metaphysics of perception.

## Overview of Block A: Enactivism and representationalism

(§2, §3 and §4)

- **Sub-question A:** *To what extent has enactivism challenged or influenced representationalism?*
- **Conclusion A1:** Perception is guided by both action and minimal representations.
- **Conclusion A2:** Enactivism must be rehabilitated to include minimal representations.

## **Block B: Enactivism and egocentricity**

### **5. Piercing the soap bubble**

#### **5.1. What is egocentricity?**

##### **5.1.1. From enactivism to egocentricity**

In §2, §3 and §4, I presented and assessed the enactivist approach to perception. Ultimately, I concluded that perception requires minimal representations, which fits closely with Clark's account, and with a rehabilitated version of O'Regan and Noë's sensorimotor enactivism. Enactivists wish to reject representations. However, we cannot fully eliminate representations, since perception still requires partial, personal, action-oriented models.

What do we mean by 'personal' and 'action-oriented'? These terms relate to the physical body of the organism and its physical place in the environment. The organism must represent relevant actions (expectations arising from sensorimotor contingencies) and relevant environmental features (an index for the world). Partial, personal, action-oriented models of the world are relevant to the organism and relative to the organism. I will argue that this should be understood in terms of egocentricity. I will use the term 'egocentric' to capture the sense in which a representation is relative to the capabilities and needs of an organism. Egocentricity is a highly useful way to reframe the notions of action and representation. When we say

that a representation is action-oriented, we mean that it relates, in an egocentric way, to the bodily capabilities and needs of an organism. This is the first step along the path to understanding the place of the perceiver in perception – a path that will culminate in metaphysics, in §6 and §7.

### **enactivism + egocentricity**

(pairwise interaction)

- When we say that a representation is action-oriented, we mean that it relates, in an egocentric way, to the bodily capacities and needs of an organism.
- Thus, sensorimotor contingencies and action-oriented models should be understood in terms of egocentricity.

Egocentricity is also a highly useful way to understand and analyse enactivism. The notion of action is at the heart of enactivism, and this is best understood in terms of egocentricity. Moreover, as we will see, O'Regan and Noë's sensorimotor contingencies should be understood as egocentric representations. They are representations which relate, fundamentally, to the bodily abilities of the organism. When Tiger Woods, master of movement, perceives the first fairway, he represents his possibilities for action: *if I swing the driver with this velocity, then I will perceive the ball whizz off in that direction*. These are Tiger's egocentric expectations about his own potential actions.

Here, in §5, I will discuss egocentricity and its counterpart, allocentricity. I will investigate the extent to which perception is shaped by egocentricity. And I will show that egocentricity is crucial for understanding enactivism: sensorimotor contingencies and the place of the perceiver.

### **5.1.2. The world inside a soap bubble**

Every creature on Earth lives in its own 'soap bubble'.<sup>93</sup> Jakob von Uexküll offers this poetic articulation of the inherent egocentricity of perception. In his view, every organism – whether an ant, a bee, a chimpanzee or a human – inhabits a separate representational world. Every organism inhabits a different visual space. Each of these worlds is defined by the sensorimotor capacities of the organism, and each is populated with different stimuli and possibilities for action. These are the separate 'soap bubbles' that enclose every organism.<sup>94</sup>

How small are these soap bubbles? According to von Uexküll, the visual space of a human infant extends to a radius of only ten metres, compared to between six and eight kilometres for a human adult. After this 'farthest plane', it becomes impossible to distinguish the true size and shape of objects – whether they are small or simply far away.<sup>95</sup> Such a comparison merely highlights the differences within the human species, let alone the different visual spaces of insects, birds and other mammals.

---

<sup>93</sup> Von Uexküll, 'A stroll through the worlds of animals and men', 5.

<sup>94</sup> Von Uexküll, 'A stroll through the worlds of animals and men', 28.

<sup>95</sup> Von Uexküll, 'A stroll through the worlds of animals and men', 26–7.



These different ‘worlds’ are also characterised by differences of possibility. This is best explained in terms of ‘affordances’, a notion at the centre of Gibson’s ecological psychology. The idea is that organisms perceive objects differently, because objects *afford* different possibilities to different organisms.<sup>96</sup> A human and a bee will not perceive a flower in the same way, since a flower presents a human and a bee with different possibilities for action. When any organism interacts with the world, it perceives only a narrow subset of the stimuli present in the environment, constructing a perceptual world that is relevant for its survival. That is von Uexküll’s core claim about separate ‘soap bubbles’.

### **5.1.3. Piercing the soap bubble**

Von Uexküll’s elucidation of soap bubbles was first published in 1934 so, in the meantime, almost a century has passed. Yet von Uexküll is still a formidable opponent. For the purposes of this thesis project, he makes an ideal opponent, since he articulates the precise type of egocentricity and anti-realism that I am striving to overcome. The essay in question, ‘A stroll through the worlds of animals and men’, is written with a popular audience in mind, a feature which brings more benefits than drawbacks. So, while it is lacking in technical details, Von Uexküll provides an arsenal of beautiful quotations and lyrical imagery. This makes him all the more engaging as an opponent. In this section, I will return to von Uexküll at various points, responding

---

<sup>96</sup> Gibson, *The ecological approach to visual perception*, 127.

to his anti-realist articulations. I will also bring his notions in line with the contemporary, neuroscientific context.

The terms ‘egocentric’ and ‘allocentric’ will be a core focus of this section, and I will clarify these terms in much greater detail in the next subsection (§5.2). In brief, egocentricity refers to a self-centred view of the environment, whereas allocentricity refers to an other-centred or world-centred view. The notion of egocentricity encapsulates von Uexküll’s core argument about soap bubbles. Such separate representational worlds are egocentric, since they are constructed based on stimuli relevant to the subject. In fact, von Uexküll makes the case for profound egocentricity in perception. He writes: ‘There is no space independent of subjects’.<sup>97</sup> There is no space outside of individual soap bubbles. Instead, the world consists of countless egocentric spaces, which may intersect to a greater or lesser degree.

The response, in favour of allocentricity, must argue that we can *pierce* the soap bubble, in some way. To achieve allocentricity, a human must be able to look beyond her own, narrow perceptual world. She must be able to step outside of herself, to some degree. This thesis project strives to *pierce* the soap bubble, metaphorically speaking. In more technical terms, I aim to show that egocentricity does not preclude the independence of the world, and does not preclude humans’ ability to perceive features of a mind-independent world. When I talk about allocentricity, I am referring to representations of the mind-independent world – or representations of mind-independent properties or entities. Allocentric representations must represent the world as it exists outside of the subject’s perspective. Are such representations

---

<sup>97</sup> Von Uexküll, ‘A stroll through the worlds of animals and men’, 29.

possible for human beings? In §5, I will investigate the extent and the limit of allocentricity in human perception. I will formulate an answer to **Sub-question B: *To what extent should human perception be understood as ‘egocentric’ or ‘allocentric’?***

In §5.1, I have introduced von Uexküll as an opponent, and I have outlined the challenge of *piercing* the soap bubble as a way to affirm humans’ ability to represent mind-independent properties. Next, I will discuss spatial navigation using the notions of egocentricity and allocentricity (§5.2). These terms – egocentricity and allocentricity – can be fruitfully applied to issues in the philosophy of perception, and I will show that they encapsulate core claims from Gibson, Noë and other positions linked to enactivism (§5.3). Turning to consider the notion of *perspective*, I will argue that there is evidence for allocentricity in human perception, enabling human beings to pierce the soap bubble (§5.4). In the last section, I will finish with some closing comments (§5.5).

## Overview of Block B: Enactivism and egocentricity

(§5)

- **Sub-question B:** *To what extent should human perception be understood as 'egocentric' or 'allocentric'?*
- **Conclusion B1:** Perception is guided by both egocentric and allocentric elements.
- **Conclusion B2:** Enactivism should be understood in terms of egocentricity and allocentricity.

## 5.2. Egocentricity and allocentricity in spatial navigation

### 5.2.1. Defining egocentric and allocentric

In §5.2 and §5.3, I will focus on human spatial navigation, drawing on the philosophical and neuroscientific literatures. In the context of spatial navigation, what is the best way to frame the egocentric/allocentric distinction? As a rough starting place, let's consider *what* is represented: the self (egocentric) versus entities other than the self (allocentric). Desirée Colombo and her co-authors define egocentricity and allocentricity as follows. The egocentric frame involves information about the position of the navigating self in the environment – in terms of 'subject-to-object relations' and

‘body-centred representations’. The allocentric frame involves information about the positions of objects – in terms of ‘object-to-object relations’ and ‘world-centred representations’.<sup>98</sup> There is a distinction between *what* is represented: positions and relations relative to the self (egocentric) versus positions and relations relative to other objects (allocentric). It is also important to note that ‘egocentric’, in this context, does not carry any normative implications. I am using ‘egocentric’ as a technical term only, and not in the sense of being morally self-centred.

Is there more to this difference between egocentric and allocentric representations? In July 2021, I conducted a two-hour interview with Professor Pete Mandik for my podcast, *Extrapolator*, and I had the opportunity to probe him further.<sup>99</sup> For Mandik, the difference cannot be fully reduced to *what* is represented: the self (egocentric) versus other entities in the world (allocentric). You have to say something about *how* it is represented or *how* it is referred to. For allocentricity, the mode of representation matters. The representational scheme must have a high degree of generality, in a way that abstracts away from the self.<sup>100</sup>

For example, consider how the gas gauge in an automobile represents the fuel level. A gas gauge doesn’t say: ‘the fuel in vehicle #1782 is low’ (allocentric). It just says: ‘this car is low on gas’ (egocentric).<sup>101</sup> It’s just a needle pointing closer to the icon of the gas can. It’s representing in an indexical way – in a pointing way. It can only represent this vehicle, because it’s only installed in this vehicle. It couldn’t represent

---

<sup>98</sup> Colombo et al., ‘Egocentric and allocentric spatial reference frames in aging’, 605.

<sup>99</sup> Allen (host), ‘Pete Mandik: Perceiving the world’ (podcast audio, forthcoming).

<sup>100</sup> Allen (host), ‘Pete Mandik: Perceiving the world’ (podcast audio, forthcoming).

<sup>101</sup> Allen (host), ‘Pete Mandik: Perceiving the world’ (podcast audio, forthcoming).

any arbitrarily selected vehicle. However, if I use a linguistic scheme, like a vehicle identification number (#1782), I'm able to represent it in a way that *abstracts* from the particular context of use of the gas gauge of that particular vehicle, and – as Mandik attests – now we're getting closer to allocentric territory.<sup>102</sup> Mandik's example should be qualified by the recognition that the label 'vehicle #1782', alone, is not sufficient for allocentricity. However, it has elements that are amenable to allocentricity. If we compile a catalogue or reference book with other numbered cars, we can start to refer to the fuel in vehicle #1782 in an allocentric way: in a way that abstracts away from the vehicle, rather than pointing indexically back towards the vehicle.

Thus, for Mandik, the defining feature of an egocentric representation is not just *what* is represented but *how* it is represented. An egocentric representation represents environmental relations relative to the self (*what*), in a way that is indexical and points to the self (*how*). By contrast, an allocentric representation represents entities other than the self (*what*), in a way that abstracts away from the self and points away from the self (*how*). That is a crucial takeaway for the upcoming discussion. When we are investigating the extent and the limit of allocentricity in human perception, we must pay attention not only to *what* is represented but *how* it is represented. In order to achieve allocentricity in perception, humans must be able to represent other objects in a way that abstracts away from the self – in a way that points away from the self. Allocentricity entails that level of generality and abstraction, outside the subject's perspective, outside the subject's soap bubble.

---

<sup>102</sup> Allen (host), 'Pete Mandik: Perceiving the world' (podcast audio, forthcoming).

### 5.2.2. Linking allocentricity and mind-independence

Introducing the notions of egocentric and allocentric provides a crucial link to the empirical, neuroscience literature. Egocentric representations are associated with activity in the posterior parietal cortex, whereas allocentric ones are associated with the hippocampus.<sup>103</sup> So, when we enter a philosophical discussion of perception in terms of egocentricity and allocentricity, we can – at the same time – refer to empirical studies on perception and the brain. I will be referring to neuroscientific evidence at various intervals, especially in connection to allocentricity and activity in the hippocampus.

There is also a conceptual link between egocentric/allocentric and ‘mind-dependent’/‘mind-independent’. These terms connect logically with the definitions I have already given. If a representation is allocentric, it necessarily refers to entities which persist outside the mind of the subject, and it necessarily refers to those entities in a way that abstracts away from the subject. By contrast, a representation that is egocentric depends on the mind of the subject, since it involves the subject describing himself and pointing back to himself. Thus, I will, at times, talk about mind-dependence and mind-independence in the context of egocentricity and allocentricity.

Although there is a logical connection between allocentricity and mind-independence, this connection is by no means uncontroversial. Anti-realists, in particular, would mount objections – for instance, that the definition itself begs the question of realism and mind-independence.<sup>104</sup> I will consider this objection in detail,

---

<sup>103</sup> Mandik, ‘The neural accomplishment of objectivity’.

<sup>104</sup> We have not yet entered the discussion of metaphysics (§6 and §7), but a metaphysical objection is already lurking. When I define allocentric representations in terms of mind-independent objects and

in §7.3.2. I am ultimately trying to connect the perception debate with issues in metaphysics, so I will need to justify this link to mind-independence. Realism and mind-independence will be a central focus in §7.

### **5.2.3. Egocentricity: A world with me at the centre**

Mandik outlines egocentricity with respect to receptive fields (sensory neurons) and effective fields (motor neurons). Receptive fields map out the area that an organism can perceive, whereas effective fields map out the area that an organism can act upon. Retinocentric representations involve receptive fields, coding information in an egocentric way. Retinal ganglion cells and neurons each have circular fields, and the locations of these fields are defined relative to a particular cell in a particular retinal location. So, the firing of a retinal ganglion cell represents ‘the location of a stimulus in a region of retinocentric space’.<sup>105</sup> The firing of one of these cells encodes information in a way that refers to a retinal location – in a way that is, by definition, relative to the subject.

Egocentricity also comes into play with effective fields. A motor neuron has influence over a particular region of space, and this can be defined in relation to the subject. For example, an individual’s arms are a certain length and can be extended to reach a limited region of space around the body, and this information is encoded in

---

relations, this might be *begging the question* of the mind-independent world. This thesis project is aimed at affirming the mind-independent world, so it is not a transgression to *presuppose* the mind-independent world in my premises and definitions? This objection of circularity (begging the question) is valid and deserves our attention, so I will address it in §7.2.3.

<sup>105</sup> Mandik, ‘The neural accomplishment of objectivity’.



‘eye-centred coordinates’.<sup>106</sup> For the average human being, with no visual impairments, the effective reach of arms is represented relative to the position of the eyes, in an egocentric way.

Temperature and thermoreceptors are an example of nonspatial egocentric representations. Representations of temperature encode ‘hazard’ or ‘harmlessness’, and details such as ‘too hot’, ‘too cold’ or ‘just right’. Obviously, these details are relative to the subject, since ‘too hot’ immediately raises the question, ‘too hot for whom?’ So, our representations of temperature are inherently egocentric.<sup>107</sup> We do not represent some absolute or allocentric measure (like ‘42.2 degrees Celsius’).

#### **5.2.4. Allocentricity: Abstracting away from myself**

After giving several examples of egocentricity, Mandik turns to allocentricity. There are indeed cases where we abstract away from ourselves, representing features of the environment in a way that does not relate to the self. These allocentric thoughts have a ‘detached or objective character’.<sup>108</sup> Mandik, of course, aims to find the neural basis for these detached thoughts, which abstract away from the subject. He points to empirical studies which focus on the hippocampus.

Classic experiments investigate spatial navigation in the rat brain, studying the effects of lesions to the hippocampus, by comparing the performance of rats both with and without such lesions. In one such study, performed by Howard Eichenbaum,

---

<sup>106</sup> Mandik, ‘The neural accomplishment of objectivity’.

<sup>107</sup> Mandik, ‘The neural accomplishment of objectivity’.

<sup>108</sup> Mandik, ‘The neural accomplishment of objectivity’.

Caroline Stewart and Richard Morris, rats without lesions outperformed rats with lesions. The ‘hippocampal damaged rats’ took much longer to reach the end of the water maze, and in some cases never reached the end goal.<sup>109</sup> These results suggest that the hippocampus plays a role in spatial navigation, and therefore allocentric representations enable an organism to navigate the environment.<sup>110</sup> These results do not prove that navigation is guided by allocentricity (and the hippocampus) alone; they simply prove that allocentric representations play some role in spatial navigation. This is promising evidence that spatial navigation requires *more* than egocentricity.

How far does allocentricity go? I have just demonstrated that allocentricity (associated with the hippocampus) plays *some* role in navigation. Is it possible to have ‘pure’ allocentricity? In other words, is it possible to represent the world in a way that makes no reference to the self – with zero indexicality, and with total generality? There are reasons to doubt this possibility. Rick Grush argues that full-blown allocentricity would be useless. A purely world-centric model, with no *index* to locate the self in that world, would serve no function to the organism. Such a model would be useless for navigating, guiding behaviour, or providing any other evolutionary advantage.<sup>111</sup> Mandik adds a supporting anecdote. The representation, ‘there is a pot boiling over on the stove’, is useless without some egocentric information – the location of the self, with respect to this pot that is boiling over. It matters greatly whether the pot is in the neighbouring room or halfway around the world.<sup>112</sup>

---

<sup>109</sup> Mandik, ‘The neural accomplishment of objectivity’.

<sup>110</sup> Mandik, ‘The neural accomplishment of objectivity’.

<sup>111</sup> Grush, ‘Self, world and space’, 88.

<sup>112</sup> Mandik, ‘The neural accomplishment of objectivity’.

Perceivers must always represent some index! We have heard this argument before, in §3.4.5, where I cited Wallis and Wright. They criticise O'Regan and Noë's notion of the world as an external memory, and write that such an external model 'provides no comfort to someone who has lost her keys'.<sup>113</sup> In other words, an allocentric model of the world is useless, unless you have some index as to *where you are* in that world, and – more importantly – where your keys are. So, we can never fully eliminate the egocentric element. There must always be some egocentric trace – some minimal index. In other words, there must always be some reference to self – some indexicality. That is the limit of allocentricity.

There is further evidence, from neuroscience, that egocentricity cannot be fully eliminated. We have already observed that lesions to the hippocampus (associated with allocentricity) impair navigational ability. Equally, lesions to the posterior parietal cortex (associated with egocentricity) impair navigational ability. This demonstrates that the 'hippocampus is not the sole locus of navigational ability'.<sup>114</sup> The insight here is very similar to the one about allocentricity and the hippocampus. The evidence supports the fact that egocentricity (and the posterior parietal cortex) plays *some* role in spatial navigation: that spatial navigation requires more than allocentricity (and the hippocampus) alone.

So, it appears that *both* the hippocampus *and* the posterior parietal cortex are important for navigation – *both* allocentricity *and* egocentricity. Going forwards, I will

---

<sup>113</sup> Wallis and Wright, 'Enactivism's vision', 259.

<sup>114</sup> Mandik, 'The neural accomplishment of objectivity'.

consider what it means for human navigation to be guided by a mixture of egocentric representations and allocentric representations.

### 5.2.5. Elementary versus complex spatial representations

Tobias Meilinger and Gottfried Vosgerau analyse human spatial navigation. They recognise that humans are able to employ both egocentric and allocentric representations of spatial locations, and humans sometimes prefer one or the other strategy, depending on the situation.<sup>115</sup> What does it mean to ‘prefer’ an egocentric navigational strategy versus an allocentric navigational strategy? It simply means that humans have certain tendencies and dispositions, statistically speaking, when it comes to spatial navigation tasks. Elsewhere in the literature, it has been reported that individuals ‘prefer’ to construct models of spatial relations by working from left to right.<sup>116</sup> Further, when objects are listed in the same sentence, individuals ‘prefer’ to locate these objects adjacent to one another in a spatial model.<sup>117</sup> In §5.2.6, I will discuss ‘preferences’ with respect to egocentric and allocentric navigation strategies.

The preference for egocentric versus allocentric strategies varies, depending on the task, but – often – a mixed approach is useful. For example, in the spatial task, ‘*recognizing scenes*’, subjects use both egocentric relations (a viewpoint from experience) and allocentric relations (from background objects).<sup>118</sup> Overall, Meilinger

---

<sup>115</sup> Meilinger and Vosgerau, ‘Putting egocentric and allocentric into perspective’, 207.

<sup>116</sup> Jahn, Knauff and Johnson-Laird, ‘Preferred mental models in reasoning about spatial relations’, 2075.

<sup>117</sup> *Ibid.*, 2075.

<sup>118</sup> Meilinger and Vosgerau, ‘Putting egocentric and allocentric into perspective’, 208.

and Vosgerau wish to argue that egocentric and allocentric elements frequently interact. Past research has focused on distinctions, but future research should focus on the interaction between egocentric and allocentric reference frames.<sup>119</sup>

Meilinger and Vosgerau argue that ‘elementary’ representations are insufficient for many navigational tasks.<sup>120</sup> Elementary representations render only *one* kind of relation (egocentric *or* allocentric), from *one* point of reference (or no point of reference). Elementary representations are insufficient for representing larger spaces. For example, it would be very inefficient to represent a whole city using an egocentric-only reference frame, and the amount of information involved would surely exceed the computational limits of the human brain.<sup>121</sup> This inefficiency becomes clear, once we think about the computational requirements of mapping a whole city in this way. Let’s imagine a purely egocentric map of Utrecht, where everything is represented by reference to me (in an indexical way). No object-to-object relations are allowed. I might start to list relations:

I am 4.2km from the Dom Tower.

I am 4.5km from Janskerkhof.

It takes me seventeen minutes to walk to the nearest Albert Heijn.

This ‘elementary’ representation of Utrecht is highly inefficient and, arguably, useless. It contains no information about the relation of the Dom Tower to Janskerkhof. For

---

<sup>119</sup> Meilinger and Vosgerau, ‘Putting egocentric and allocentric into perspective’, 218.

<sup>120</sup> Meilinger and Vosgerau, ‘Putting egocentric and allocentric into perspective’, 208–9.

<sup>121</sup> Meilinger and Vosgerau, ‘Putting egocentric and allocentric into perspective’, 213.

navigating large spaces, we need some strategy other than using ‘me’ as the point of origin.

Thus, to represent larger spaces, more complex forms of representations are required – those which mix egocentric and allocentric relations. Meilinger and Vosgerau discuss the need for ‘complex spatial representations’. Even a relatively simple spatial example, like ‘the target is left of the tree’, involves three locations – speaker, target and tree – enlisting *both* egocentric *and* allocentric relations.<sup>122</sup> This representation contains an allocentric relation (the target is beside the tree) *and* an egocentric relation (the target appears to be *left* of the tree, from my vantage point). Even in simple and commonplace representations, we mix egocentric and allocentric relations.

Meilinger and Vosgerau consider how a system could mix egocentric and allocentric reference frames, by using an ‘egocentric hierarchy’. Lower-level coordinates systems (both egocentric and allocentric) are all ‘subsumed under a top-level egocentric reference frame’.<sup>123</sup> Thus, in the earlier example with the target and the tree, the allocentric relations involving the tree ‘inherits its orientation’ from the top-level egocentric frame.<sup>124</sup> In this way, lower-level allocentric relations (between a tree and a rock) can all be incorporated into an overarching coordinate system, with the self at the centre. This allows the self to navigate the world, recognising relations between objects in the world, and still maintaining an overall conception of how these miniature coordinate systems relate to the self.

---

<sup>122</sup> Meilinger and Vosgerau, ‘Putting egocentric and allocentric into perspective’, 214.

<sup>123</sup> Meilinger and Vosgerau, ‘Putting egocentric and allocentric into perspective’, 216.

<sup>124</sup> Meilinger and Vosgerau, ‘Putting egocentric and allocentric into perspective’, 216–7.

Meilinger and Vosgerau's argument supports the view that egocentric coordinates play an essential role in navigation at all times. Their proposed egocentric hierarchy allows for allocentric representations, but all within the wider view of the self's place in the physical environment. This bolsters our earlier conclusion, from §5.2.5, that navigation requires both egocentric and allocentric elements. As argued by Mandik and by Wallis and Wright, the self's location in the environment can never be completely erased from perception. At the very least, humans need an egocentric marker, to indicate how far you are from the pot that's boiling over on the stove, or to indicate where you've put your missing car keys. That is the minimal *index*: the location of the self in the world. Thus, we can state **Conclusion B1: Perception is guided by both egocentric and allocentric elements**. I will briefly discuss what this means for individuals navigating the environment.

#### 5.2.6. Variation in human spatial abilities

Colombo and her team conduct a systematic review of the research on allocentric and egocentric frames in ageing. They, too, report that successful navigation does not involve one frame – egocentric or allocentric – in isolation. Rather, successful navigation requires a flexible approach, switching and combining different spatial strategies.<sup>125</sup> This finding supports **Conclusion B1** that egocentricity can never be fully erased from perception. Any achievement of allocentricity (or, per Mandik,

---

<sup>125</sup> Colombo et al., 'Egocentric and allocentric spatial reference frames in aging', 605.

‘accomplishment of objectivity’) will only ever be *partial* – with some egocentricity retained.

Colombo and her team report some other findings which have interesting implications. The overarching result from research on ageing is that spatial abilities do not remain consistent throughout life; they ‘develop from childhood to adulthood’ and then ‘deteriorate with aging’.<sup>126</sup> This means that infant humans and elderly humans do not have the same navigational abilities as humans in the middle of life. These effects appear to concern allocentric frames. An ‘allocentric decline’ can be observed from about 70 years old. Older adults are much more likely to adopt an egocentric strategy in spatial navigation tasks, compared to a younger control group, even when an egocentric strategy has no clear advantage.<sup>127</sup> So, when it comes to spatial navigation abilities, there is a variability between ages.

This allocentric decline was measured and reported by Jan Wiener and his team. They assessed participants on two tasks: route repetition (remembering turns when moving in the same direction of travel) and route retracing (remember turns when moving in the opposite direction of travel).<sup>128</sup> The first task is connected to egocentric strategies, since it involves remembering the route taken by the subject. The second task is connected to allocentric strategies, since it requires ‘allocentric computation of landmarks’ and abstracting away from the subject’s route.<sup>129</sup> In retracing a route, subjects must construct an allocentric representation of how the

---

<sup>126</sup> Colombo et al., ‘Egocentric and allocentric spatial reference frames in aging’, 605.

<sup>127</sup> Colombo et al., ‘Egocentric and allocentric spatial reference frames in aging’, 611.

<sup>128</sup> Wiener, Kmecova and de Condappa, ‘Route repetition and route retracing, 1.

<sup>129</sup> Colombo et al., ‘Egocentric and allocentric spatial reference frames in aging’, 611.



route would have looked from the opposite direction, from a different perspective. Wiener and his team found that elderly participants performed worse at route retracing, but when it came to route repeating there was no difference between the elderly and the young.<sup>130</sup> This is key evidence in favour of an allocentric decline in older people.

There is a further variability to note, besides the difference between older and younger people; there is also a variability between individuals in general. In fact, Colombo and her co-authors report a 'great variability' in the way that individuals 'preferentially use egocentric or allocentric strategies'.<sup>131</sup> This variability is underpinned by neurobiological differences, since 'allocentric learners' have 'more grey matter in the hippocampus compared to egocentric learners'.<sup>132</sup> What should we make of these differences in navigational abilities – between ages and between individuals more generally? The bottom line is: humans do not share the same ability to avail of allocentric strategies for navigation. Throughout life, some humans have a greater preference and ability for allocentric strategies, coupled with hippocampal differences. Then, later in life, elderly humans have reduced allocentric abilities. This means that humans do not all have the same ability to represent the world outside their perspective. All humans are tethered to some trace of egocentricity (a minimal index), but some humans are *more* bound to their egocentricity than others. By contrast, allocentric learners – those with an increased hippocampus – have an increased capacity for world-centred representations.

---

<sup>130</sup> Wiener, Kmecova and de Condappa, 'Route repetition and route retracing', 1.

<sup>131</sup> Colombo et al., 'Egocentric and allocentric spatial reference frames in aging', 605.

<sup>132</sup> Colombo et al., 'Egocentric and allocentric spatial reference frames in aging', 606.

Can we conclude that allocentric learners have a greater ability to *perceive the mind-independent world*? They certainly have an increased ability to *represent the self-independent world*. Allocentric learners *do* have a greater ability to view properties and relations in the world itself, outside of any perspective. These are empirically measured differences in human perception. However, the link between the neuroscience of perception and metaphysical realism is undeniably tenuous. Can we say, with any certainty, that allocentric strategies have anything to do with *metaphysics*? That is the link that I aim to establish, in §6 and §7. I will argue that evidence of allocentricity is sufficient to prove that humans *do* perceive invariant, mind-independent properties and relations ‘out there’ in the world.

In the meantime, in the next subsection, I will show how fruitful it is to apply the notions of egocentricity and allocentricity to the philosophy of perception (§5.3). I will unite a range of philosophical accounts under the umbrella of egocentricity, and I will do so in roughly chronological order – starting with historical authors like von Uexküll and Gibson, followed by contemporary authors like Noë. This discussion will link back to enactivism. I will show that the notions of egocentricity and allocentricity are key to understanding the claims of enactivists. This will build towards **Conclusion B2: Enactivism should be understood in terms of egocentricity and allocentricity.**

Enactivists do not use the terms egocentricity and allocentricity to discuss perception, but I will argue that that it is a highly fruitful approach. It is not that enactivists are mistaken, or that they commit some omission, by using other terms. Other explanatory strategies may be suitable. However, it is certainly the case that the

terms egocentricity and allocentricity allow us to shed new light on enactivist claims, especially with respect to the metaphysics of enactivism.

### 5.3. Applying egocentricity and allocentricity to enactivism

#### 5.3.1. Von Uexküll: The *Umwelt*

Some time has passed since von Uexküll first discussed ‘soap bubbles’, in 1934, and some time has passed since *I* last discussed them, in §5.1. Von Uexküll’s core claims are encapsulated by the notion of egocentricity. He makes the case for profound egocentricity, arguing that every organism is trapped inside its own egocentric bubble. Von Uexküll introduces his specialist term, *Umwelt* (plural: *Umwelten*), to describe an animal’s soap bubble.<sup>133</sup> The *Umwelt* is a unified egocentric world, informed by all the ways the subject can act and perceive, and all the ways it cannot. These sensorimotor abilities and needs constrain the representations that are available to an organism, and constrain the entities or stimuli that an organism attends to. Such constraints give rise to a unique *Umwelt* for each organism.

Von Uexküll illustrates this notion of an *Umwelt* by describing the world of a female tick. The tick is a tiny creature, who hangs from the foliage of bushes, seeking the hot blood of a mammal. The tick has no eyes, though her skin has a ‘general photosensitivity, so she finds her perch by climbing up towards the sunlight.’<sup>134</sup> On this ‘watchtower’, the tick waits. She is waiting for a signal – the scent of butyric acid

---

<sup>133</sup> Von Uexküll, ‘A stroll through the worlds of animals and men’, 6.

<sup>134</sup> Von Uexküll, ‘A stroll through the worlds of animals and men’, 6–7.

- which is released 'from the skin glands of all mammals'.<sup>135</sup> When the tick detects butyric acid, she drops from her branch. Her descent is haphazard, merely a blind hope that she will land on a hot, succulent mammal. When she lands, she can detect the temperature of the new surface. If it is cold, this means that the tick has not found her prey, and she starts to climb again towards the light. However, if she detects warmth, she has hit her target. The tick scurries around 'to find a hairless spot' and then burrows, drinking the mammal's blood.<sup>136</sup>

The most remarkable insight here relates to the *narrow* existence of the tick: her impoverished representational world. Of all the stimuli in the wide world - of all the stimuli available to human beings (the rainbow of visible light, not to mention other sense modalities) - the tick perceives almost nothing. She detects only three stimuli: 'the chemical stimulus' (butyric acid), 'the temperature stimulus' (skin) and 'the mechanical stimulus' (hairs).<sup>137</sup> Von Uexküll emphasises this point: 'out of the hundreds of stimuli radiating from the qualities of the mammal's body, only three become the bearers of receptor cues for the tick'.<sup>138</sup> Compared to human sensory capacities, the tick perceives almost nothing. Yet she perceives *enough* - enough to navigate her own, narrow representational world and enough to effect her own, narrow outcomes. This is the *Umwelt* of the tick: 'Out of the vast world which surrounds the tick, three stimuli shine forth from the dark like beacons, and serve to lead her unerringly to her goal'.<sup>139</sup>

---

<sup>135</sup> Von Uexküll, 'A stroll through the worlds of animals and men', 7.

<sup>136</sup> Von Uexküll, 'A stroll through the worlds of animals and men', 7.

<sup>137</sup> Von Uexküll, 'A stroll through the worlds of animals and men', 11.

<sup>138</sup> Von Uexküll, 'A stroll through the worlds of animals and men', 11.

<sup>139</sup> Von Uexküll, 'A stroll through the worlds of animals and men', 11-12.

Moreover, we would be mistaken to pity the tick with her impoverished existence. Here, 'poverty' may be a virtue rather than a vice. The 'poverty' of the tick's representational world enables and empowers the tick to attain the blood necessary for her survival. This insight becomes a maxim for von Uexküll: 'security is more important than wealth'.<sup>140</sup> In other words, perceiving a *rich* representational world, with a *wealth* of stimuli, does not entail any evolutionary advantage. If the outcome is *survival* ('security'), then an impoverished soap bubble may do just fine.

Of course, von Uexküll really wants to emphasise the claim that each organism – each tick, bee, chimpanzee or human – inhabits its own *Umwelt*. This is von Uexküll's anti-realist claim: that we each inhabit a separate, egocentric world, and that there is no mind-independent space between (or outside of) these worlds. Interestingly, he very occasionally slips into realist talk, as in the sentence: 'out of *the hundreds of stimuli radiating from the qualities of the mammal's body*, only three become the bearers of receptor cues for the tick'.<sup>141</sup> Von Uexküll appears to treat the 'hundreds of stimuli' as objective, mind-independent regularities. It does not appear to matter that the tick can perceive only three. The other hundreds of stimuli persist independently of the tick's awareness or ignorance, tacitly supporting the idea that stimuli are mind-independent states of affairs, *outside* any particular soap bubble. However, apart from this rare lapse, von Uexküll entertains an anti-realist tack.

In von Uexküll's argument about 'security' versus 'wealth', he is pitting the two constraints against one another. In his view, the primary evolutionary goal of

---

<sup>140</sup> Von Uexküll, 'A stroll through the worlds of animals and men', 12.

<sup>141</sup> Von Uexküll, 'A stroll through the worlds of animals and men', 11 (emphasis added).

survival does *not* favour rich representational worlds. In fact, it may favour the opposite. In the interests of efficiency and parsimony, each organism selectively represents those narrow stimuli necessary for its survival. There is no incentive in favour of representational wealth, von Uexküll argues. The result is a multiplicity of vastly different representational worlds – vastly different *Umwelten*. Where, then, is the shared, mind-independent world? That is the challenge that we face, in responding to von Uexküll. His colourful account argues against the existence of any world outside of individual soap bubbles. In §5.5, I will sketch my response to von Uexküll – my argument about how humans can *pierce* the soap bubble.

### 5.3.2. Gibson: Ecological psychology and affordances

Gibson's ecological psychology was a prominent feature in the previous sections. I discussed his theory of direct visual perception, which I argued was 'true anti-representationalism'. Gibson claims that perception is completely unmediated. For Gibson, perception involves the direct pickup of information by the organism. Gibson uses the term 'affordances' to describe what the organism directly perceives. According to his definition: 'The affordances of the environment are what it offers the animal, what it provides or furnishes, either for good or ill'.<sup>142</sup> The organism perceives the possibilities that the environment *affords*.

Gibson's theory of affordances can be explained in terms of egocentricity. If a human [for example, Tiger Woods] perceives something as too large to eat, then this

---

<sup>142</sup> Gibson, *The ecological approach to visual perception*, 127.

perception depends on facts about himself: his size, metabolism, and so on.<sup>143</sup> Tiger Woods sees the environment in terms of what *he* can do, in terms of the opportunities available to *him*. So, perceiving affordances depends on an egocentric view of the environment. Similarly, when a mouse perceives a crack in the skirting board as a timely hiding place, this depends on facts about the mouse's body – its size and sensorimotor ability to scurry across the room. The cat, who is chasing the mouse, is much too big for the hole and does not perceive it as a relevant hiding place; the crack in the skirting board does not *afford* this possibility to the cat. The crack is still the same crack, on an allocentric view, but it results in very different egocentric representations in the mouse and the cat.

I'll provide one last example, in homage to von Uexküll's lyrical style. Imagine a wild meadow, teeming with flowers and insects in the heat of a summer afternoon. This particular meadow is located deep in the Dutch countryside, in the southern province of Limburg, only a few hours' walk from the city of Maastricht. Here, the low, rolling hills present some rare elevation in an otherwise flat landscape. The meadow is soaked in warm sunshine. Long tangled grasses cover the hard, cracked earth.

Now imagine a bee and a human (Bethany), located at one end, looking out over the meadow. Of course, there will be a difference in colour perception; bees, unlike humans, can detect ultraviolet light, whereas bees cannot detect orange and red. So, the human and the bee will each see different 'colours', so to speak. In *Figure*

---

<sup>143</sup> Mandik, 'The neural accomplishment of objectivity'.

5, the bee's soap bubble is shifted towards the blue end of the spectrum, depicting the difference in light detection.

Moreover, the human and the bee will represent vastly different possibilities for action. For the bee, it will be possibilities to pollenate, as markings on flowers – known as nectar guides – are made more visible due to ultraviolet detection. In the soap bubble of the bee, ultraviolet nectar guides are highly relevant stimuli which draw the bee's attention, as shown by small red arrows in *Figure 5*. For Bethany, the human, the meadow may be a small blip in a wider journey. Bethany may look past the twenty metres of tangled wildflowers and trample ahead for a few short seconds, to reach the ice-cream van parked on the road. On an allocentric view, it is the same meadow, with the same flowers, insects and automobiles. However, it gives rise to vastly different egocentric representations in Bethany and the bee, based on the differences in what the environment *affords* to each of them.





Figure 5: Bethany and the bee

Here, the notion of egocentricity neatly sums up the core of Gibson's argument. He wants to argue that an organism perceives only those features of the world that are relevant for its survival or success. The crucial point here is that perception is guided by those properties or relations which afford possibilities to the organism. What the environment *affords* to the organism is clearly determined by egocentric facts about the organism. Consider the earlier example of the cat and the mouse, and the notion of hiding places:

An important kind of place, made intelligible by the ecological approach to visual perception, is a place that affords concealment, a *hiding place*. The concealing of oneself from other observers and the hiding of a detached object from other observers have different kinds of motivation. As every child discovers, a good hiding place for one's body is not necessarily a good hiding place for treasure.<sup>144</sup>

For Gibson, what we perceive is guided by the possibilities that are afforded. It matters greatly whether a mouse is trying to hide its own body, or whether a cat is trying to hide its own body, or whether a cat is trying to hide a dead bird. In each case, what counts as a hiding place is constrained by bodily needs and bodily capacities. In other words, it is constrained by egocentric facts about the organism. Equally, for Bethany and the bee, what counts as 'relevant' stimuli in the meadow is constrained by bodily needs and bodily capacities. Perception is guided by egocentric constraints.

### 5.3.3. Noë: Sensorimotor enactivism

Noë discusses Gibson's theory of affordances and reformulates it to fit the enactive approach. For Noë, seeing involves perceiving the relevant sensorimotor contingencies. This means that we don't see retinal images of objects, or representations of objects. What's more, we don't see the raw perceptual objects themselves; we see possibilities for movement.<sup>145</sup> When we 'see' a steep hill or a meadow full of flowers, we are perceiving the possibilities for action that the hill or

---

<sup>144</sup> Gibson, *The ecological approach to visual perception*, 136.

<sup>145</sup> Noë, *Action in perception*, 105–6.

meadow affords to us. Thus, Gibson's theory of affordances can be brought into line with enactivism. An organism's sensorimotor profile defines its possible actions, and defines what possibilities the environment affords. For Noë, affordances are determined by sensorimotor contingencies.

Meilinger and Vosgerau make a further, exciting link between enactivism and the notion of egocentricity. It is exciting not only for the purposes of this thesis project, but also for shedding further light on the nature of sensorimotor contingencies. Representing sensorimotor contingencies involves representing the 'contingencies between action and perception' – how performing certain bodily actions will produce a change in sensory input.<sup>146</sup> Mastery of sensorimotor contingencies involves *knowing* how sensory input will change when performing these actions. Crucially, sensorimotor contingencies are – by definition – relative to the self. Masters of movement know what sensory changes *their own actions* and *their own body* will produce. Thus, *sensorimotor contingencies are egocentric representations*. They are 'truly egocentric', because space is represented in terms of possible actions, relative to the subject's sensorimotor capacities.<sup>147</sup>

In Noë's own words, mastery involves a 'thoughtless automaticity' around our bodily abilities: 'We spontaneously crane our necks, peer, squint, reach for our glasses, or draw near to get a better look (or better to handle, sniff, lick or listen to what interests us)'.<sup>148</sup> Of course, on correct assessment, these are *egocentric* ways of engaging with the world. Our peering, squinting and sniffing depends on our bodily abilities.

---

<sup>146</sup> Meilinger and Vosgerau, 'Putting egocentric and allocentric into perspective', 209.

<sup>147</sup> Meilinger and Vosgerau, 'Putting egocentric and allocentric into perspective', 209.

<sup>148</sup> Noë, *Action in perception*, 1–2.

So, when we represent the world in terms of sensorimotor contingencies, we are representing the various ways that we might interact with the world, within an egocentric frame of reference. Thus, we arrive at the key insight of **Conclusion B2: Enactivism should be understood in terms of egocentricity and allocentricity**. We should conceive of sensorimotor contingencies as egocentric, since they relate to bodily abilities and needs.

**Conclusion B2** is centrally important for this thesis project, since it allows us to connect the three main elements: enactivism, egocentricity and realism. Since enactivism can be understood in terms of egocentricity, we can now assess the metaphysical claims of enactivism in terms of egocentricity. In §6 and §7, I will show that the most central issues in the metaphysics of enactivism can be expressed in terms of egocentricity and realism. I will also show how egocentricity and realism can be reconciled – in the spirit of symmetry, and in the context of the enactivist's account of perception.

In §5.4, I'll consider one last puzzle involving egocentricity and allocentricity. So far, in §5, I have been discussing human spatial navigation. Now, in §5.4, I will enter the literature on object recognition. I will investigate the extent to which object recognition is guided by egocentric or allocentric elements. In particular, I will explore the phenomenon of *perspective*, in the context of object recognition. Take a human subject, Bethany. Is she trapped within her own perspective (egocentric)? Or, can she abstract away from her perspective (allocentric)? Can she perceive only perspectival shapes (egocentric), or can she infer the distal (allocentric) shapes of objects?

## 5.4. Perspective and object recognition

### 5.4.1. The dual aspect of perceptual content

Noë presents a very interesting discussion of perspectival content. He argues that perceptual content has a ‘dual aspect’.

Perceptual content has a dual aspect. There’s the way experience presents the world as being... apart from your perspective... And there is *the way* the world is presented in experience... from your vantage point.<sup>149</sup>

In Noë’s view, experience is both object-centred and self-centred. If Tiger Woods holds his favourite golf club horizontally and looks along its length, like a marksman looking down the barrel of a gun, he perceives it *both* in an object-centred way (as a 110cm rod) *and* in a self-centred way (as little more than a black dot, with most of the length occluded from view).

We can rephrase this ‘dual aspect’ in terms of egocentricity and allocentricity, even though Noë does not rely on these terms. Perspectival content is both egocentric (self-centred, from some vantage point) and allocentric (object-centred, irrespective of vantage point). Noë gives the example of a plate. When you see a plate from an oblique angle, it looks circular *and* elliptical.<sup>150</sup> When Tiger Woods looks across the table at his friend’s empty dinner plate, he can perceive *both* the object-centred shape of the plate

---

<sup>149</sup> Noë, *Action in perception*, 163.

<sup>150</sup> Noë, *Action in perception*, 163–4.

(circular) *and* the self-centred shape of the plate (elliptical). The plate really looks circular, and it really looks elliptical from his vantage point.

There is a crucial link between perspectival content and enactivism. Noë claims that enactivism can account for the dual nature of content. The plate looks elliptical (egocentric aspect) because perception is grounded in action; perception is defined by facts about the organism's body and vantage point.<sup>151</sup> The fact of perspective 'marks the place of action in perception'.<sup>152</sup> The plate looks circular (allocentric aspect) because knowledge of sensorimotor contingencies allows objects to be reconstructed as they *are*.<sup>153</sup> We can say that the allocentric shape of the plate is circular, because that is its true, objective shape, irrespective of any vantage point. When Tiger Woods grasps the sensorimotor profile of a plate, he understands its allocentric shape, and how changes in vantage point can cause changes in egocentric appearance. Tiger Woods, master of crockery, knows that an elliptical dinner plate is simply a circular plate, viewed from an oblique angle. In this way, enactivism explains *both* aspects of content – egocentricity because of action and the physical body, allocentricity because of sensorimotor knowledge.

Here, we find further support for **Conclusion B2: Enactivism should be understood in terms of egocentricity and allocentricity**. Although Noë doesn't use these terms, he is trying to express the 'dual nature' of perspectival content in terms of egocentric and allocentric elements. Moreover, he is trying to explain the core claims of enactivism in terms of this dual nature. On the one hand, perception is

---

<sup>151</sup> Noë, *Action in perception*, 164.

<sup>152</sup> Noë, *Action in perception*, 34.

<sup>153</sup> Noë, *Action in perception*, 164.

constrained by action, the physical body and the physical location in the environment (egocentric). On the other hand, through sensorimotor knowledge, we can understand the invariant nature of relations and properties outside of our own perspective (allocentric).

It is important to note that perspective in object recognition is just one subset of the phenomena that can be analysed in terms of egocentricity and allocentricity. Earlier in §5, I discussed spatial navigation, which is another subset. None of these phenomena is perfectly equivalent to the egocentricity/allocentricity distinction. Indeed, as we saw with spatial navigation, human often (and, perhaps, necessarily) employ a mixture of strategies. So, these are umbrella terms, which provide a tool for analysis. My ongoing claim, **Conclusion B2**, is that egocentricity and allocentricity are a fruitful tool for understanding enactivism.

#### 5.4.2. Perspectival properties

It is useful to introduce the terms ‘P-property’, ‘P-shape’ and ‘P-size’. Noë employs these terms in his discussion of perspectival content, as a shorthand for ‘perspectival property’, ‘perspectival shape’ and ‘perspectival size’.<sup>154</sup> Shape (allocentric) is distinct from perspectival shape or ‘P-shape’ (egocentric). Tiger Woods can perceive both the circular shape and the elliptical P-shape of the empty dinner plate. Equally, size is distinct from perspectival size or ‘P-size’. When Tiger Woods looks across the golf

---

<sup>154</sup> Noë, *Action in perception*, 83.

course and sees a small copse of trees, the P-size of the trees is only a few centimetres – no bigger than his thumb – yet he knows that the trees are several metres in height.

Amusing explanations of perspective can be found in pop culture. *Father Ted* is a classic Irish sitcom from the 1990s about a group of priests living on a remote island. The priests are chaotic and incompetent, getting into a series of mishaps – a topic which was still controversial in Catholic Ireland at the time. The episode ‘Hell’ features an iconic scene where Father Ted attempts to explain perspective to Father Dougal, who is characteristically dim-witted. Ted holds up two tiny toy cows and says: ‘these are *small*’. Then, he points out the window to a group of real cows in a distant field and says: ‘but the ones out there are *far away*’. Dougal looks intensely puzzled, shaking his head. ‘*Small... far away,*’ repeats Ted, by way of summary. Dougal continues shaking his head and Ted throws the toy cows down on the table, exclaiming, ‘ah, forget it’.<sup>155</sup>

---

<sup>155</sup> The 30-second clip is available at: <https://www.youtube.com/watch?v=MMiKyfd6hA0> (Accessed 1 April 2021).





Figure 6: Father Dougal gets confused about perspectival properties

Father Dougal has not mastered the relevant laws of sensorimotor contingency, so he is confused about perspectival properties. He does not understand why the P-size of the tiny toy cows is the same as the P-size of the real cows in the distant field. The scene is amusing precisely because *no* adult routinely makes such a mistake. Aside from the occasion visual illusion, we are masters of perceiving a world of distant objects, easily seeing size (allocentric) as distinct from P-size (egocentric). In fact, Noë argues that we pay little attention to P-properties in normal life.<sup>156</sup> We look beyond apparently tiny cows and tiny trees and elliptical shapes. We look from afar and straightforwardly perceive regular-sized cows and regular-sized trees and perfectly circular plates. Such is our mastery of sensorimotor knowledge.

---

<sup>156</sup> Noë, *Action in perception*, 83.

What is the metaphysical status of P-properties? Moreover, if there are multiple simultaneous viewers of the same object, does this bring about multiple P-properties? The short answer is: yes. If five humans are viewing the same cow from different distances, that cow will have a different P-size for each of the observers.

Let's imagine a cow in a field, and then pepper the scene with five people. Observer 1 (Tiger Woods) is only two metres from the cow and he's enthusiastic to have such a clear view of the cow's muscular rear. Observer 2 (Queen Elizabeth II) is eleven metres from the cow, though she insists on measuring the distance as thirty-six feet. Her majesty's enthusiasm has also been aroused by the sight of the handsome cow and his glossy, brown coat. Observer 3 (Nicole Kidman) and Observer 4 (Barack Obama) stand fifty and eighty metres from the cow, respectively. Finally, Observer 5 (poor Father Dougal), is 150 metres away, in the neighbouring field. At that distance, his confusion about perspectival properties is unlikely to be alleviated.

For each of the five observers, the cow appears to have a different size. Up close, to Tiger Woods, the cow looks almost as tall as its true, allocentric height of two metres. To poor Father Dougal, the cow appears no different in size to the toy cow that he clutches in his hand. For Father Dougal, the perspectival size of the cow in the distant field is the same length as his finger. In his hand, he holds a toy cow which is also the same length as his finger. Such are the puzzling delights of perspective.

Metaphysically, each of these P-properties arises from the *relation* between cow and observer. Perspectival size is not a property contained in the cow itself, unlike its allocentric size. Allocentric size is a physical property of the cow, whereas perspectival size is a property of the observer's *relation* to the cow; perspectival size only arises due

to geometric facts about an observer's vantage point with respect to the cow. So, multiple P-properties persist at the same time, and they are not metaphysically independent of the relation between observer and cow. P-properties are egocentric in nature, so they arise only *because of* such relations, and they are subject to change, following changes in those relations. In §7, I will present a detailed discussion of perspectival properties, and I will show how they are compatible with metaphysical realism.

Here, the most important implication is *how much* allocentricity is enjoyed by human perceivers. Thanks to our mastery of sensorimotor knowledge, all the observers (except for Father Dougal) can abstract the true, allocentric size of the cow, as distinct from its mere perspectival size. In our everyday perceptions, we regularly abstract the invariant properties of entities like distant cows and trees. In other words, we are not fooled by cows and trees that appear small. We routinely observe distant cows and understand the cow's true size (allocentric), even though it appears small from our vantage point (egocentric). This further bolsters **Conclusion B1: Perception is guided by both egocentric and allocentric elements**. In the case of object recognition, just as in the case of spatial navigation, there is a dual aspect to perception.

### 5.4.3. The debate about perspectival shape

I will conclude this section with some empirical research on perspective and object recognition. Jorge Morales and his co-authors take an empirical approach, investigating perspectival shape and distal shape in perception. They use the example

of a rotated coin. A coin's true, distal shape is circular, though a rotated coin appears elliptical, from a certain perspective. Such a rotated coin gives rise to an elliptical image on the retina: 'it projects an ellipse on the back of the eye'.<sup>157</sup>

The question, for perception research, is whether we "see" the coin as circular or elliptical (or both). Morales and his team view these as the three possible states of affairs:

1. Option 1: We "see" the elliptical shape of the rotated coin, and then infer that it is, in fact, a circular coin. This was the view held by British empiricists, like Locke and Hume – that we perceive perspectival properties.<sup>158</sup>
2. Option 2: We "see" the circular shape of the coin, even when it is rotated, and it takes effort and reflection to realise that the rotated coin is projecting an elliptical shape. This was the view held by Gibson and Helmholtz – that 'we primarily perceive 3D distal properties'.<sup>159</sup>
3. Option 3: We "see" the coin as *both* elliptical (perspectival shape) *and* circular (distal shape) at the same time.<sup>160</sup> This is, of course, the view espoused by Noë. He writes: 'Perceptual content has a dual aspect. There's the way experience presents the world as being... apart from your perspective [allocentric]... And there is *the way* the world is presented in experience... from your vantage point [egocentric]'.<sup>161</sup>

---

<sup>157</sup> Morales et al., 'Sustained representation of perspectival shape', 14873.

<sup>158</sup> Morales et al., 'Sustained representation of perspectival shape', 14873.

<sup>159</sup> Morales et al., 'Sustained representation of perspectival shape', 14873.

<sup>160</sup> Morales et al., 'Sustained representation of perspectival shape', 14873.

<sup>161</sup> Noë, *Action in perception*, 163.

Morales and his co-authors write that Option 2 is the orthodox view in contemporary vision science. They use the terms ‘the orthodox view’ and ‘the mainstream view’.<sup>162</sup> The orthodoxy holds that perception tracks distal properties, not perspectival properties.<sup>163</sup> In other words, perception is guided by object-centred representations, rather than self-centred representations.<sup>164</sup> Let’s rephrase these positions in terms of egocentricity and allocentricity.

1. Option 1: Egocentricity guides perception and action. We perceive perspectival properties.
2. Option 2: Allocentricity guides perception and action. We perceive distal properties.
3. Option 3: Perception has a dual aspect, guided by *both* egocentricity *and* allocentricity. We perceive (or detect) *both* perspectival properties *and* distal properties.

The orthodoxy (Option 2) holds that allocentricity guides perception and action. Morales and his team set out to investigate the empirical basis for the orthodox position.

---

<sup>162</sup> Morales et al., ‘Sustained representation of perspectival shape’, 14881.

<sup>163</sup> Morales et al., ‘Sustained representation of perspectival shape’, 14873.

<sup>164</sup> Morales et al., ‘Sustained representation of perspectival shape’, 14881.

#### 5.4.4. Experiments and findings

Morales and his co-authors conduct nine experiments to empirically test the role of perspectival properties in perception. Subjects had to search for a distally elliptical coin, and avoid being distracted by distally circular coins which had been rotated (so as to appear elliptical). Morales and his team hypothesised that subjects would ‘be “distracted” by a rotated circle’, if its projection matched the distal ellipse beside it. This distraction would lead to a delay in response time. So, if subjects took longer to identify the true, distal ellipse, this would be evidence that they had been ‘distracted’ by the rotated coin, proving that perspectival shape plays a role in perception.<sup>165</sup> *Figure 7* and *Figure 8* are images taken from the paper by Morales and co-author, showing the difference between a rotated circle (which appears elliptical) and a distal ellipse (which is truly elliptical).<sup>166</sup>



*Figure 7: A rotated circular coin*

---

<sup>165</sup> Morales et al., ‘Sustained representation of perspectival shape’, 14874.

<sup>166</sup> Morales et al., ‘Sustained representation of perspectival shape’, 14874–5.



Figure 8: A distal ellipse (left); a rotated circle (right)

Indeed, in all nine experiments, subjects were distracted by the perspectival shape of rotated coins. They exhibited slower response times when a perspectival ellipse matched the distal ellipse of their target. Morales and his co-authors conclude that: 'An elliptical coin is harder to distinguish from a rotated circular coin... because the two objects *appear to have something in common*'.<sup>167</sup> This demonstrates that perspectival shapes play an influential role in perception. The perspectival ellipse of the rotated coin is not tacit or irrelevant; rather, it guides perception and action, causing distraction and slower response times.

These results challenge the orthodox view (Option 2) that allocentricity guides perception. These results also echo the conclusion that we have encountered throughout §5: that we cannot completely erase egocentricity from perception. Perspectival properties – information about the place of the *self* in the representation – are always retained. That is the minimal *index* – the 'I am here' on the allocentric

---

<sup>167</sup> Morales et al., 'Sustained representation of perspectival shape', 14880.

map – that is retained, even in the most permissive case. So, egocentricity and the fact of perspective cannot be eliminated from perception.

Perhaps the most interesting conclusion from Morales and his co-authors states that ‘objects have a remarkably persistent dual character: their objective shape “out there,” and their perspectival shape “from here”’.<sup>168</sup> Here, we find experimental evidence for Noë’s claim that experience has a dual aspect: both self-centred and object-centred.<sup>169</sup> Our perceptions of the world are still stamped with our egocentric perspective. However, at the same time, we can abstract the allocentric properties of distal objects. The two persist alongside one another. This is a final piece of empirical evidence in favour of **Conclusion B1: Perception is guided by both egocentric and allocentric elements.**

## 5.5. Closing comments on egocentricity and allocentricity

### 5.5.1. The dual nature of human perception

We are now in a position to formulate an answer to our second sub-question: *to what extent should human perception be understood as ‘egocentric’ or ‘allocentric’?* In §5, I have presented empirical and philosophical evidence for **Conclusion B1: Perception is guided by both egocentric and allocentric elements.** Both elements are present in human perception, as strategies for navigation, or as aspects of perceptual content.

---

<sup>168</sup> Morales et al., ‘Sustained representation of perspectival shape’, 14873.

<sup>169</sup> Noë, *Action in perception*, 163–4.



How far does allocentricity go? When it comes to spatial navigation, humans can employ allocentric strategies, representing environmental relations in a way that abstracts from their own perspective. These strategies are evidenced by activity in the hippocampus, and they enable a human to take an object-centred view of the environment. When it comes to perceiving objects, a human can detect allocentric properties in spite of her perspective. A human can identify the allocentric shape of a circular coin, even though it looks elliptical from her vantage point.

However, allocentricity has a limit; that is the other side of the 'coin'. Egocentricity can never be fully erased from perception. When it comes to spatial navigation, a human must always retain an egocentric index – which says, 'I am here'. When it comes to perceiving objects, there is empirical evidence that perspectival properties play a role in guiding perception. Egocentricity is associated with slower response times when identifying objects. Thus, we are left with **Conclusion B1**, an account of human perception that includes *both* egocentric *and* allocentric elements.

The discussion of egocentricity and allocentricity also has important implications for enactivism. Sensorimotor contingencies should be understood as egocentric representations, because they catalogue a perceiver's egocentric ways of engaging with the environment, based on bodily capacities and needs. Furthermore, enactivism (à la Noë) comfortably accounts for the dual nature of perception, both egocentricity and allocentricity. The egocentric aspect of perception arises from the role of action, the physical body and the physical location in the environment. Perception is constrained by egocentric facts about the perceiver. These constraints also inform the sensorimotor contingencies which connect the perceiver's perceptions

and potential actions. Yet, on the other hand, enactivism accounts for our routine perception of allocentric properties and relations. We can abstract allocentric shapes (like the distal circle of a coin) and allocentric spatial relations, through sensorimotor knowledge. If we understand how transformations in our vantage point will alter perspectival shape, then we can abstract how the distal shape of the object persists, apart from any perspective. Enactivism provides a robust and plausible explanation for both aspects of this dual nature of perception. Moreover, the enactivist's account of perception aligns neatly with the notions of egocentricity and allocentricity. We can use these terms to consolidate the enactivist's account. Thus, I argued for **Conclusion B2: Enactivism should be understood in terms of egocentricity and allocentricity.**

The link between enactivism and egocentricity is crucial for this thesis project. It enables us to discuss the metaphysics of enactivism in terms of egocentricity and realism. In §6 and §7, the three main elements in this thesis project – enactivism, egocentricity and realism – will come together, in my critical discussion of the metaphysics of enactivism.

### **5.5.2. The prospect of piercing the soap bubble**

We return, at last, to von Uexküll's soap bubbles. We have, in a way, vindicated von Uexküll's claims. Organisms in general, and humans in particular, *do* rely on egocentric strategies to perceive the world, whether for spatial navigation or for identifying objects. However, I have also presented evidence against von Uexküll's more radical claims. It is not egocentricity all the way down, as von Uexküll imagines. The empirical evidence supports the notion of allocentricity in perception, and this

allows humans to *pierce* the soap bubble. There are strong arguments in favour of allocentric strategies for spatial navigation, aided by the hippocampus, and in favour of the claim that humans can identify allocentric properties, like shape and size. Thus, there is evidence for allocentricity in perception – evidence that human beings *can* perceive invariant properties ‘out there’ in the world. Human beings are *not* trapped within a purely egocentric soap bubble. We can pierce through.

I conclude that human beings *can* pierce the soap bubble, which means that the aim of Block B is complete. Yet we are immediately greeted by the core questions for the next section. What does piercing the soap bubble mean for our ability to perceive the metaphysical world? What is the link between allocentricity and reality ‘out there’? What is the link between perception and metaphysics? These questions will populate the remaining sections, §6 and §7. In the rest of this thesis project, I will investigate the link between metaphysical reality and human perception, outlining my overall account of the metaphysics of enactivism.

## Overview of Block B: Enactivism and egocentricity

(§5)

- **Sub-question B:** *To what extent should human perception be understood as 'egocentric' or 'allocentric'?*
- **Conclusion B1:** Perception is guided by both egocentric and allocentric elements.
- **Conclusion B2:** Enactivism should be understood in terms of egocentricity and allocentricity.

## Block C: Enactivism and metaphysics

### 6. The matrix: Representationalism and realism

#### 6.1. Where representationalism and realism meet

##### 6.1.1. Entering the realism debate

In Block C, at last, we enter the realism debate. I will formulate an answer to **Sub-question C: *What is the link between enactivism and realism, as claimed by other authors and as justified by the evidence?*** Throughout the previous sections, I have been dropping signposts and cliffhangers for the discussion ahead. The main focus of this thesis project is the *metaphysics* of enactivism, and I have already alluded to the metaphysical implications of O'Regan and Noë's enactivism, Varela, Thompson and Rosch's (onto-epistemological) anti-foundationalism and Gibson's claims about optic-array information. More intriguing still, I pointed to the *agnosticism* of Clark – the view that anti-representationalism entails *nothing*, in either direction, about realism or anti-realism.

Block C is divided into two parts. In §6, I map out the connections between the representationalism and realism debates. I describe this first part as 'procedural', since it is a fairly mechanical exercise in sketching out the range of positions. It is only later that I engage, critically, with the substance of the arguments. Nonetheless, I argue that the procedural task is hugely valuable. It is also my first major input – an original contribution to the literature, I believe. No other philosopher, as far as I am aware, has

attempted such an extensive *survey* or *synthesis* of the connections between the representationalism and realism debates.

I describe the second part, §7, as ‘substantive’. Here, I engage with the substance of the arguments for and against realism. Since this thesis project is focused primarily on the enactivist approach, I pay special attention to enactivism, and evaluate the salience of the *metaphysics* of enactivism. We stand to learn a lot from enactivism, especially from the way it characterises the role of the perceiver. However, some enactivists (O’Regan and Noë) make tacit metaphysical presumptions. The fact that these presumptions are unsubstantiated is a weakness in their account, and there is a need for an explicit defence of realism, in enactivist terms. Moreover, since other enactivists (Varela, Thompson and Rosch) are explicitly anti-realist, there is an even greater need for an account that affirms metaphysical realism.

Throughout this thesis project, I have been engaging with the enactivist approach to perception, and it has been a wonderful guide to the most crucial issues. So much of enactivism is careful, innovative and astute. However, like any young research programme, it comes with its flaws. Since I have taken a particular interest in the *metaphysics* of perception, I stand to critique enactivism from a unique vantage point. O’Regan and Noë’s own account is incredibly light on metaphysics, which leads to concerns about their tacit claims. With the proper focus on metaphysics, we can set the record straight. I will continue to ‘rehabilitate’ enactivism. I have already argued that O’Regan and Noë must embrace minimal representationalism in order to make their account work. Equally, they must incorporate an explicit defence of realism, since it is another element that their account depends upon.

## Overview of Block C: Enactivism and metaphysics

(§6, §7)

- **Sub-question C:** *What is the link between enactivism and realism, as claimed by other authors and as justified by the evidence?*
- **Conclusion C1:** Enactivists disagree about metaphysical implications, claiming realism, anti-realism or agnosticism.
- **Conclusion C2:** Enactivism (and, in particular, the egocentric profile of a human perceiver) is compatible with metaphysical realism.
  - **The argument from biology:** Metaphysical realism is compatible with egocentric facts about embodiment.
  - **The argument from geometry:** Metaphysical realism is compatible with egocentric facts about embeddedness.

### 6.1.2. Defining 'realism' and 'anti-realism'

This thesis project is concerned with metaphysical realism about perceptual objects. In this context, realism involves the claim that the objects we perceive *exist*, and that they are *independent of* our representations (or other features of subjectivity). Drew Khlentzos provides some precise definitions of metaphysical realism:

Metaphysical realism is the thesis that the objects, properties and relations the world contains, collectively: the structure of the world... exists independently of our thoughts about it or our perceptions of it.<sup>170</sup>

According to metaphysical realism, the world is as it is independent of how humans or other inquiring agents take it to be. The objects the world contains, together with their properties and the relations they enter into, fix the world's nature and these objects [together with the properties they have and the relations they enter into] exist independently of our ability to discover they do.<sup>171</sup>

The two most important features in this definition are (i) existence and (ii) independence. Thus, metaphysical anti-realism is opposed to one (or both) of these features:

Anti-realists either doubt or deny the existence of the structure the metaphysical realist believes in or else doubt or deny its independence from our conceptions of it.<sup>172</sup>

In the context of perception and enactivism, the issue of 'independence' is particularly relevant. As we will see, those who are anti-realist about perceptual objects – von Uexküll, Chemero, and Varela, Thompson and Rosch – argue that the world is *not* independent of our representations. For von Uexküll, the only 'world' that exists is the perceptual world represented by each organism, and that world is *not*

---

<sup>170</sup> Khlentzos, 'Challenges to Metaphysical Realism'.

<sup>171</sup> Khlentzos, 'Challenges to Metaphysical Realism'.

<sup>172</sup> Khlentzos, 'Challenges to Metaphysical Realism'.



independent of representations. Therefore, von Uexküll is a metaphysical anti-realist (see §6.3.4). Varela, Thompson and Rosch make quite a different argument. They argue that perceptual objects are brought into being (are ‘enacted’) through the interaction between perceiver and environment. So, even though they reject representations, perceptual objects still are not independent of the actions and interactions of the perceiver (see §6.3.3). That is a brief flavour of the anti-realist positions that we will encounter, and why the issue of independence is particularly relevant.

In §7, I will argue in favour of metaphysical realism. I will argue that objects, properties and relations – the entities subject to human perception – exist independently of representations (and other features of the perceptual process). Specifically, I will show that enactivism is compatible with metaphysical realism. This becomes my overall case for the metaphysics of enactivism. By appealing to enactivism, we can give a robust account of human perception, accounting for both the place of the perceiver (‘the egocentric profile’) and the persistence of the world (metaphysical realism).

## enactivism + realism

(pairwise interaction)

- Anti-realists argue that enactivism entails anti-realism, because the world is not independent of the perceiver; rather, the perceiver *enacts* perceptual objects through dynamical interactions with the environment.
- Realists argue that enactivism is compatible with realism, because the world is still independent of the perceiver.

## 6.2. The *matrix*

### 6.2.1. Connecting *representationalism* and *realism*

Figure 9 is a visual map of the relations between the representationalism debate and the realism debate. It takes the form of a 3x3 matrix. Looking along the *y*-axis, you will see that I have allocated three possible positions on representationalism:

(-) Anti-representationalism

(0) Neutral/middle ground

(+) Pro-representationalism

Equally, looking along the *x*-axis, you will notice three possible positions on realism:

(-) Anti-realism

(0) Neutral/middle ground

(+) Pro-realism

Thus, my 3x3 matrix depicts a grid of nine possible positions, each of which is a *combination* of: a position about representationalism *and* a position about realism. The four corners of the grid show four major philosophical traditions:

(-, +) Idealism

(+, +) Cognitivism

(-, -) Anti-foundationalism

(+, -) Direct realism

In the context of enactivism, the term ‘anti-foundationalism’ has a slightly non-standard meaning. As I outlined in §3.3, Varela, Thompson and Rosch argue for a distinct type of ‘onto-epistemological’ anti-foundationalism.<sup>173</sup> Their anti-foundationalism goes beyond the mere question of foundations in human knowledge, à la Otto Neurath.<sup>174</sup> Instead, they also reject the idea of metaphysical foundations. So, when I use the term ‘anti-foundationalism’, I am referring to Varela, Thompson and Rosch’s far-reaching onto-epistemological proposal. I will analyse (and refute) their brand of anti-realism in §7.3.3.

---

<sup>173</sup> Vörös and Riegler, ‘A plea for not watering down the unseemly’, 3.

<sup>174</sup> Neurath, ‘Anti-Spengler’, 199.

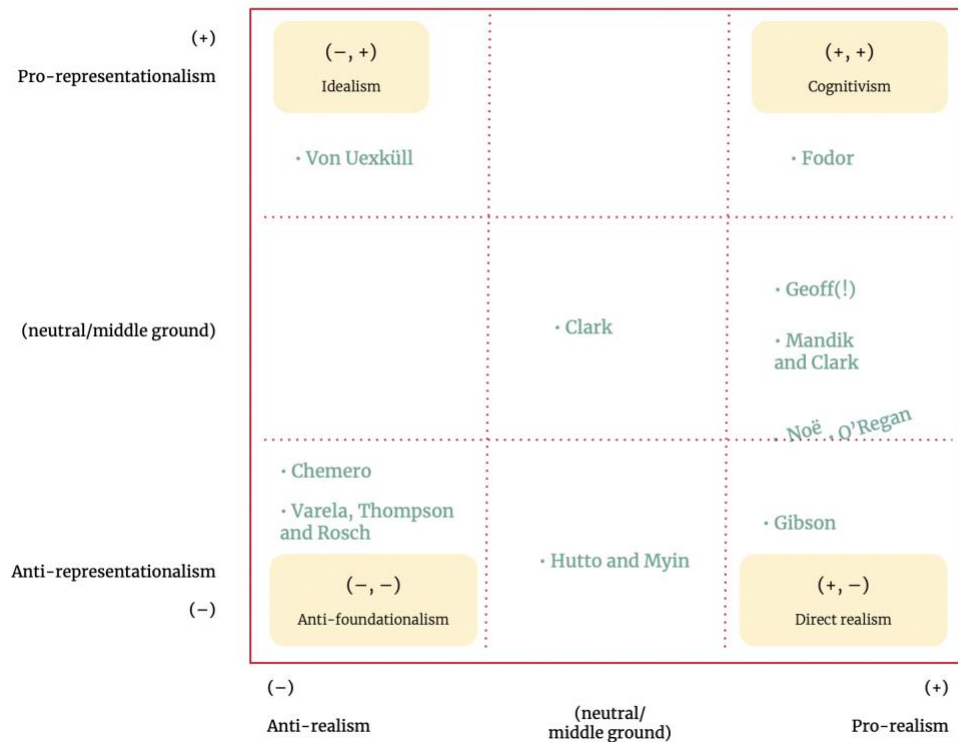


Figure 9: A 3x3 matrix comparing positions on representationalism with positions on realism

When I began researching the various forms of enactivism – from Varela, Thompson and Rosch, to Noë and O'Regan, to Hutto and Myin – one of my earliest insights was with respect to the variety of metaphysical positions. Thus, I set out to answer **Sub-question C: What is the link between enactivism and realism, as claimed by other authors and as justified by the evidence?** As I will outline in §6.3, enactivists make a great range of claims about the links to realism. Varela, Thompson and Rosch are explicitly anti-realist. O'Regan and Noë are implicitly realist. Hutto and Myin are explicitly agnostic – arguing that enactivism entails nothing about metaphysical realism. I will analyse the salience of each of these responses: realist, anti-realist and agnostic.

I have already testified that no philosopher has published such an extensive survey, as the one shown at *Figure 9*. However, a handful of secondary sources make some (narrow and minor) connections between enactivism and realism. For instance, Joel Parthemore considers Noë's metaphysics. While Noë assumes realism, 'much of the enactive community' assumes *anti*-realism. On Parthemore's own assessment, the enactivist's 'emphasis on the inseparability of agent and environment' more logically supports anti-realism.<sup>175</sup> Here, Parthemore is following Varela, Thompson and Rosch's onto-epistemological anti-foundationalism, shown in the bottom-left corner of *Figure 9*. Their extreme idea of inseparability points towards anti-realism, since neither object nor perceiver exist independently; their existence is co-defined.

Parthemore also makes a link between cognitivism and realism, more specifically, against 'Fodor's *informational atomism* theory, which stands or falls on its brand of realism'.<sup>176</sup> This was another invigorating insight, near the start of my thesis project, that some *pro*-representationalist accounts assume realism (like Fodor's informational atomism) and some *anti*-representationalist accounts assume realism (like Noë's enactivism). Throughout §6.3, I will be amassing evidence in favour of **Conclusion C1: Enactivists disagree about metaphysical implications, claiming realism, anti-realism or agnosticism**. It is not the case that there is a one-dimensional relation between realism and enactivism. Instead, realism can be accepted or rejected by both camps. This insight marks the importance of the procedural task – mapping out the grid of comparison (*Figure 9*). It also marks the importance of interrogating all

---

<sup>175</sup> Parthemore, 'From a sensorimotor to a sensorimotor++ account of embodied conceptual cognition', 145.

<sup>176</sup> *Ibid.*, 145.

possible positions and to determine which is most salient. Such a wide matrix of accepted positions demonstrates widespread confusion (or at least widespread disagreement) – and a greater need to delineate the metaphysics of enactivism.

In the next section (§6.3), I will ‘plot the points’ on the grid, taking each philosopher or doctrine one at a time and locating them in the matrix: Gibson (§6.3.1), O’Regan and Noë (§6.3.2), Varela, Thompson and Rosch (§6.3.3), Von Uexküll (§6.3.4), Hutto and Myin (§6.3.5) and Clark (§6.3.6). The primary focus of this thesis project is enactivism, which is an anti-representationalist position, so the bottom half of my grid is much more heavily populated. In future work, by me or by other philosophers, one could focus on the top half of the grid, populating it with the various positions on idealism and cognitivism. That work is beyond the scope of this thesis project, though it is important and interesting work for future projects.

### **6.2.2. The explicit and the tacit**

The last thing to mention, in this preamble, is whether the metaphysical stance of a philosopher or doctrine is ‘explicit’ or ‘tacit’. Many authors do take an explicit stance as to the metaphysics of perception. Gibson argues in favour of realism, while Varela, Thompson and Rosch argue against realism. However, it was even more important to catalogue the philosophers who make *tacit* assumptions and presuppositions about realism. Most crucially, O’Regan and Noë almost never mention metaphysics, and yet their form of enactivism is parasitic on direct realism. This is good news and bad news. The bad news is: O’Regan and Noë’s account suffers from a metaphysical lacuna. The good news is: they create honest work for lowly, young philosophers. I thank

professors O'Regan and Noë for providing such gainful employment during the past nine months.

Author	Realism?	Explicit or tacit?
Gibson	(+)	explicit
Mandik and Clark	(+)	explicit
Chemero	(-)	explicit
Varela, Thompson and Rosch	(-)	explicit
Clark	0	explicit
Hutto and Myin	0	explicit
Noë	(+)	tacit
O'Regan	(+)	tacit
Von Uexküll	(-)	tacit

*Table 3: Comparing 'explicit' versus 'tacit' positions on realism*

Are O'Regan and Noë required to provide an explicit account of the metaphysics of sensorimotor enactivism? Or, are they licenced to remain tacit on metaphysics? In §6.3.2, I argue that they must provide an explicit metaphysical account; they are self-professed metaphysicians, making claims about perception, mind and world – claims with much metaphysical import. I hope that this project is helpful in providing an explicit account of the metaphysics of enactivism, which has been lacking thus far.

### 6.3. Plotting the points

#### 6.3.1. Gibson

First up is Gibson, an anti-representationalist and pro-realist. In his paper, 'New reasons for realism', Gibson opposes the sense-datum theory of perception, putting forward his own theory of direct perception. He argues that his way of reconceptualising human perception also adds weight to the arguments in favour of direct realism. On my assessment, Gibson is certainly correct that his theory of direct perception fits into a neat package with direct realism. However, as I will outline later, there is a problem of circularity. The theory of direct perception does not appear to be an appropriate *premise* in favour of direct realism, since direct perception already *presupposes* direct realism.

Gibson objects to 'existing theories of perception'. Writing in 1967, he was referring to the sense-datum theory, which held that perception is built on sensations and 'some kind of operation' is required to 'convert' these sensations into percepts. In opposition, Gibson puts forward 'the information-based theory of perception'. He argues that sensory impressions are not entailed in perception; they are 'incidental' rather than primitive. The organism directly picks up information from the environment, without the use of sensations or any other mediating entity.<sup>177</sup> Of course, Gibson's theory of direct perception and direct pickup should already be familiar from our earlier discussion (§3.2).

---

<sup>177</sup> Gibson, 'New reasons for realism', 162.



Gibson stresses throughout that the relevant ‘information’ consists of ‘invariants’:

The first assumption of this theory of perception is that certain properties of the energy flux at the skin of an active animal do not change, whereas other properties do. The former are invariant, the latter variant.<sup>178</sup>

Through trial and error, the organism learns to distinguish the variants from the invariants. For Gibson, this process of disambiguation is part of learning to navigate the world and happens in childhood.<sup>179</sup> Invariants are properties belonging to the environment, which do not vary under perspective transformations.<sup>180</sup> Interestingly, Gibson’s definition of invariance comes very close to my notion of allocentricity. Gibson is making an astute insight when he describes perception as the process of extracting allocentric properties – a description I would broadly agree with. However, as I outlined in §3.1, there are problems with his brand of anti-representationalism. So, instead, my account of the metaphysics of perception relies on minimal representationalism.

---

<sup>178</sup> Gibson, ‘New reasons for realism’, 163.

<sup>179</sup> Gibson, ‘New reasons for realism’, 165–6.

<sup>180</sup> Gibson, ‘New reasons for realism’, 163.

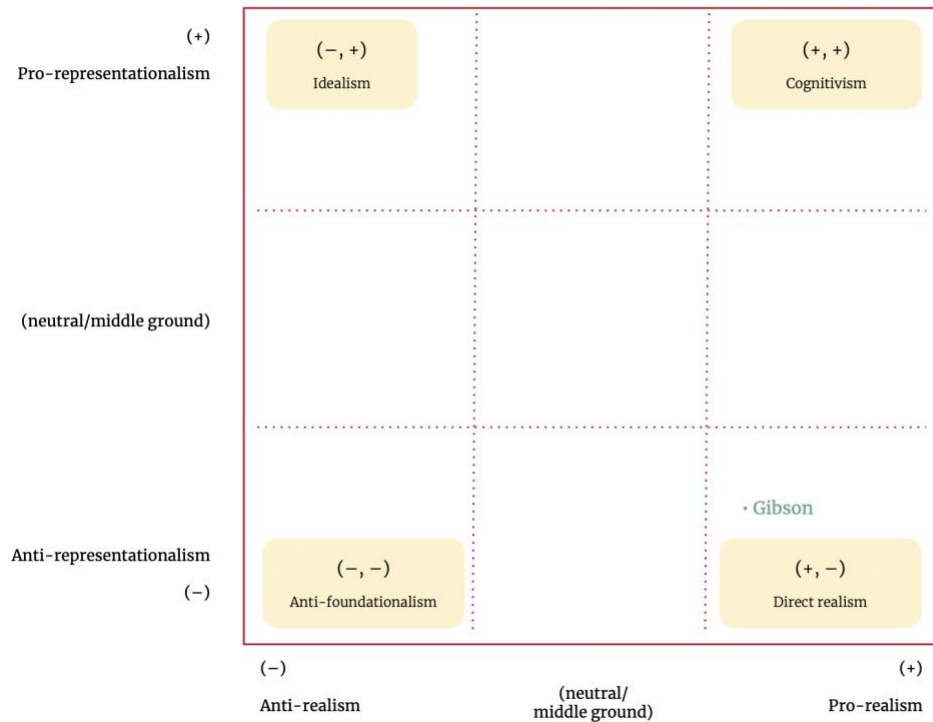


Figure 10: Gibson is *anti-representationalist* and *pro-realist*

In, 'New reasons for realism', Gibson gives an explicit defence of direct realism. He argues that his theory of direct perception makes direct realism 'reasonable'. Direct realism is a 'common-sense position', a 'naïve belief in the world of objects and events', and yet Gibson claims that he has found 'sophisticated support' for such a view. Direct realism is hard to defend on the basis on the sense-datum theory, but easy to defend when the perceptual system is considered in the context of identifying invariances.<sup>181</sup>

The most significant implication of Gibson's argument is the claim that we have 'direct or immediate awareness of objects and events'.<sup>182</sup> For Gibson, perception is

<sup>181</sup> Gibson, 'New reasons for realism', 168.

<sup>182</sup> Gibson, 'New reasons for realism', 168.

completely *unmediated*; we have direct experience of objects themselves, not of representations of objects or retinal images of objects. True ‘indirect’ perception is only possible through human artefacts – seeing a printed image or hearing a needle scratching a record. Gibson draws a forceful disanalogy to human perception:

...I am quite certain that there is no such thing as a phonograph record in the ear and no such thing as a picture in the eye – no reproduction of an external event or object that the organ transmits to the brain.<sup>183</sup>

Gibson’s theory provides the first example of how an *anti*-representationalist position can be compatible with *pro*-realism. It will be most interesting to compare Gibson with other anti-representationalists who make completely contradictory claims about realism. In §7, I will have more to say about the substance of Gibson’s metaphysics. In §7.2.3, I diagnose a vicious circularity in Gibson’s account, since he aims to prove realism, while already presupposing realism.

### 6.3.2. O’Regan and Noë

O’Regan and Noë make a tacit assumption in favour of direct realism (which I will use interchangeably with ‘naïve realism’.) This is best illustrated by the theory of the world as an outside memory, as first articulated by O’Regan, and as discussed in my earlier sections. O’Regan rejects the view that the brain constructs an internal model

---

<sup>183</sup> Gibson, ‘New reasons for realism’, 169.

of the world for the very reason that ‘it is continuously available “out there”’.<sup>184</sup> This raises the question: *is* the world straightforwardly available “out there”? O’Regan certainly squanders no time in justifying such a presumption. There is no explicit defence of realism – nor any explicit mention of realism – anywhere in his paper. The success of his theory about the external memory store rests on the truth (or salience) of naïve realism, and yet it remains tacit and unquestioned. (This allegation of ‘tacit and unquestioned’ realism is, *verbatim*, the one that Varela, Thompson and Rosch level against cognitivism.<sup>185</sup> Here, I fire it back in the direction of enactivism.)

Do O’Regan and Noë *need* to explicitly justify realism? On one view, they are working within an empirical paradigm – cognitive science – which takes realism for granted. It would be unreasonable to expect cognitive scientists to justify realism every time they publish a neurobiological study; realism is a problem for philosophers. Yet, does this work as an excuse for O’Regan and Noë? It does not. They are not mere cognitive scientists, engaged in the nuts and bolts of empirical research. They *are* philosophers – philosophers with a great deal of programmatic ambition. Moreover, they are explicit about their aims. They want to outline a ‘new metaphysics of mind and body’.<sup>186</sup> They are making claims about the *metaphysical* nature of perceivers, objects and the world, over and above physiological questions about perception. O’Regan and Noë are explicit about this goal of metaphysical reform, and yet they omit to justify a core presumption in their work – the presumption in favour

---

<sup>184</sup> O’Regan, ‘Solving the “real” mysteries of visual perception’, 464.

<sup>185</sup> Varela, Thompson and Rosch, *The embodied mind*, 9.

<sup>186</sup> O’Regan and Noë, ‘A sensorimotor account of vision and visual consciousness’, 971.

of naïve realism. Let's remember that O'Regan and Noë are philosophers (and self-professed metaphysicians), so such a presumption must be substantiated.

In their seminal paper on enactivism, 'A sensorimotor account of vision and visual consciousness', O'Regan and Noë echo and endorse the theory of the 'world as an outside memory'. O'Regan and Noë write that the outside world is 'an *external memory* that can be probed at will by the sensory apparatus'.<sup>187</sup> They also contemplate our phenomenology of vision. When we move through the world, our eyes are processing only a small number of details at any given moment, and yet we have 'the feeling of immediate availability about the whole scene'.<sup>188</sup> The whole visual field appears to be continually present, in all of its richness, even though we are only *attending to* a tiny fraction of it. O'Regan and Noë explain that this experience arises because the whole scene really *is* available to us. We must simply turn our eyes, or turn our attention, in one direction or another, and the details can be retrieved.<sup>189</sup> All the information from the visual world is stored on this external hard drive, and we are always plugged in. All the richness of the world is only one click away. Throughout O'Regan and Noë's work, this tacit commitment to realism is obvious. In §7, I will assess the substance of O'Regan and Noë's claims, and I will provide the explicit defence of realism that is lacking in their account.

I have found some, limited secondary literature on enactivism and realism. Paul Coates, in *The metaphysics of perception*, argues that enactivism presupposes direct

---

<sup>187</sup> O'Regan and Noë, 'A sensorimotor account of vision and visual consciousness', 946.

<sup>188</sup> O'Regan and Noë, 'A sensorimotor account of vision and visual consciousness', 946.

<sup>189</sup> O'Regan and Noë, 'A sensorimotor account of vision and visual consciousness', 946–7.

realism.<sup>190</sup> Here, Coates focuses on Noë's account – especially his monograph, *Action in perception* – since it is one of the most detailed and careful expositions of enactivism.<sup>191</sup> Parthemore makes a similar, albeit more general, claim: 'sensorimotor theory assumes – indeed, depends on – a particular, realist metaphysical position without making that commitment clear'.<sup>192</sup> I should note that Parthemore's argument is far less precise than Coates'. Parthemore's argument spans only a few paragraphs, whereas Coates outlines, over the course of a whole chapter, how Noë's enactivism relies on direct realism and vehicle externalism. Parthemore, at least, recognises that Noë's realist metaphysics is *implicit*. If Noë had explicitly outlined the relevant metaphysics, the uncertainty and confusion may have been avoided. If Noë had included a detailed working out of the metaphysics of sensorimotor theory, he might have pre-empted some of these criticisms. At the very least, an explicit account would have made for a sharper target.

---

<sup>190</sup> Coates, *The metaphysics of perception*, 5.

<sup>191</sup> Coates, *The metaphysics of perception*, 99.

<sup>192</sup> Parthemore, 'From a sensorimotor to a sensorimotor++ account of embodied conceptual cognition', 145.

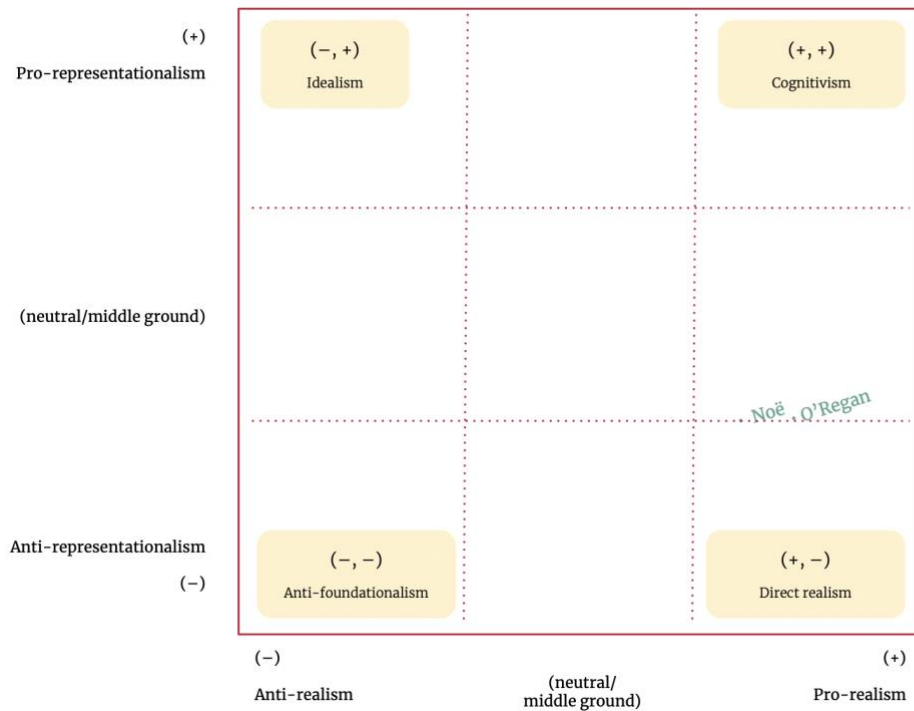


Figure 11: O'Regan and Noë are (purportedly) *anti-representationalist* and (implicitly) *pro-realist*

In Figure 11, you will notice that I have placed O'Regan and Noë *on the line* between anti-representationalism and the middle ground. Their position on representationalism is complex (and contradictory), as I discussed at length in §3. In a nutshell, they claim to be anti-representationalist, and yet sensorimotor enactivism *requires* some minimal representations on the part of perceivers. Perceivers must represent their expectations about the sensorimotor contingencies related to an action. I also argue that perceivers must represent an index – an egocentric ‘I am here’ – which locates the perceiver on any allocentric map. Thus, I defend *minimal representationalism*. Further, I argue that O'Regan and Noë must, too, be minimal representationalists. They make strong statements about rejecting representations, but their theory does not work without representations.

### 6.3.3. Varela, Thompson and Rosch

It is useful to begin this discussion of Varela, Thompson and Rosch by outlining the view that they *oppose*. In *The embodied mind*, Varela Thompson and Rosch write in opposition to ‘cognitive realism’, which refers to cognitivism and other pro-representationalist views.<sup>193</sup> These approaches might disagree about *how* representations are instantiated – symbolic and local, versus nonsymbolic and distributed – but they agree *that* cognition is representation. For Varela, Thompson and Rosch, representationalism presupposes realism. So, cognitivists and other pro-representationalists are united by a commitment to realism, hence the umbrella term, ‘cognitive realism’. In *Figure 12*, cognitivism (+, +) is located in the top-right corner, since it is pro-representationalist and pro-realist.

Varela, Thompson and Rosch are really objecting to a ‘heavy’ sense of representation, which carries epistemological and ontological commitments. This heavy sense of representations, endorsed by cognitivists, assumes that the world is *pregiven*. For cognitivists, cognition involves mentally reconstructing a world of mind-independent objects.<sup>194</sup> Varela, Thompson and Rosch object to both tenets of cognitivism – the idea that cognition *is* representation, and the idea that the world is *pregiven*. Thus, Varela, Thompson and Rosch’s account is both anti-representationalist and anti-realist.

---

<sup>193</sup> Varela, Thompson and Rosch, *The embodied mind*, 147.

<sup>194</sup> Varela, Thompson and Rosch, *The embodied mind*, 135.



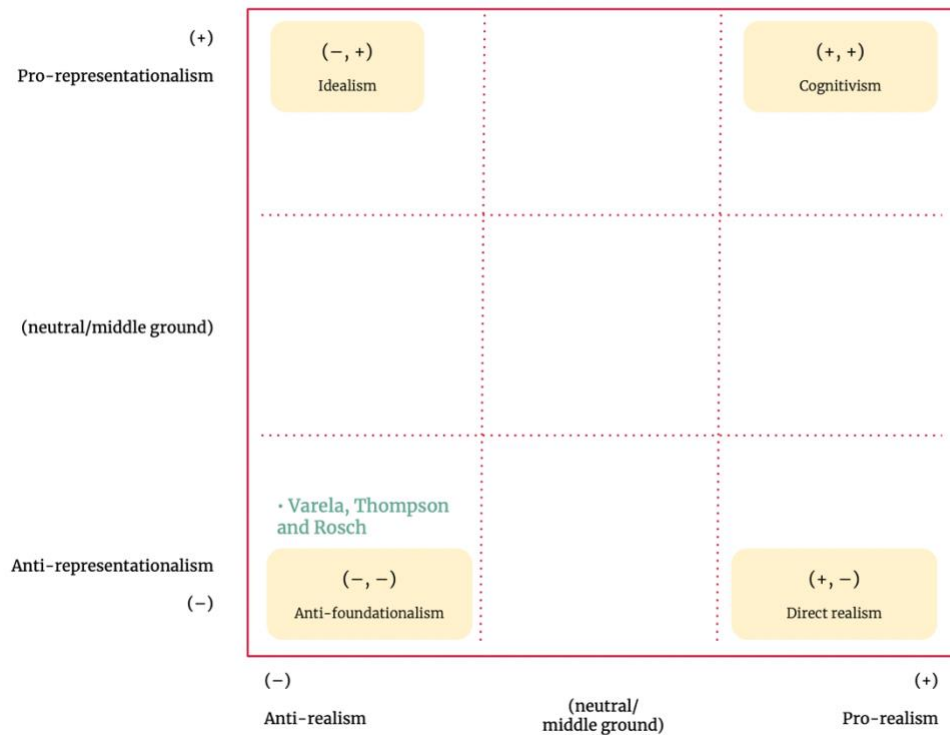


Figure 12: Varela, Thompson and Rosch are *anti-representationalist* and *anti-realist*

Varela, Thompson and Rosch's position is best described as anti-foundationalist. As outlined earlier, their anti-foundationalism is onto-epistemological, and therefore slightly non-standard. They want to reject any metaphysical foundation, whether 'the world' or 'the image'. They reject both the realist's claim to foundations (that the world comes before the image), and the idealist's claim to foundations (that the image comes before the world). It is a chicken-and-egg question, they argue. Mind and world exist in circularity, with no foundation; they specify one another.<sup>195</sup> Varela, Thompson and Rosch's views about anti-representationalism and enactivism argue for a distinct metaphysical arrangement. The dynamics of the perceiver and the dynamics of world

<sup>195</sup> Varela, Thompson and Rosch, *The embodied mind*, 172.

are wrapped up in an inherent circularity. The perceiver brings the world into being, through her interactions with the world. Perceptual objects are *enacted* by the perceiver and do not exist outside of that interaction. Thus, for Varela, Thompson and Rosch, the metaphysics of perception is one of *interaction* between perceiver and world. Outside of that interaction – outside of the feedback loop – there is no mind-independent world and there are no world-independent minds.

It is interesting to contrast Gibson with Varela, Thompson and Rosch. For Gibson, the perceiver interacts directly with the world, without any mediating representations. However, the world has a persistence that precedes cognition. The world contains variants and invariants, independent of the perceiver, and allocentric properties can be picked up in perception. Varela, Thompson and Rosch present an altogether different metaphysics. The perceiver, indeed, communes directly with the world, but the world itself is *defined* by that interaction. There is nothing invariant about the world *prior* to cognition. It is the very act of the perceiver that enacts certain features of the world.

Anthony Chemero makes a noteworthy comparison between Gibson and Varela, Thompson and Rosch. In fact, he compares the two factions of scholars that grew out of the two bodies of work. On the one hand, *ecological psychologists* (including Michael Turvey, Robert Shaw, Edward Reed and William Mace) are influenced by Gibson. On the other hand, *enactivists* (including Varela, Thompson and Rosch) are influenced by phenomenologists.<sup>196</sup> Chemero describes these two factions as ‘two

---

<sup>196</sup> Here, Chemero focuses on *The embodied mind* by Varela, Thompson and Rosch as a central opus for enactivism.

tight-knit and loosely interacting groups of scientists'.<sup>197</sup> It might be more useful to use labels for the two groups, like *Gibsonians* and *Varelans*.

How much do Gibsonians and Varelans have in common? Chemero writes: 'One of the key aims of my work has been to get ecological psychologists and enactivists to realize that they are (almost) on the same page'.<sup>198</sup> Chemero uses an umbrella term, 'radical embodied cognitive science' to unite these factions. He defines the term broadly, as 'the use of dynamical modelling to put the theoretical positions of Gibson, Heidegger and Merleau-Ponty in touch with data about perception, action and cognition that can be gathered in the lab'. He lists Varela and Thompson among the 'pioneers' of this work.<sup>199</sup> So, Chemero attempts to unite the Varelans and the Gibsonians under an umbrella term, arguing that 'they are (almost) on the same page'.<sup>200</sup> Much to the contrary, there is (literally) a world of difference between them. True Gibsonians would be anti-representationalist and *pro*-realist. Whereas, disciples of Varela and co. would be anti-representationalist and *anti*-realist. So, in fact, Chemero is gravely mistaken. Gibsonians and Varelans have completely opposing views on fundamental metaphysics.

These are the crucial, metaphysical insights that sometimes go unnoticed in the philosophy of perception. This is why the procedural task is so important. It may appear that Gibsonians and Varelans have similar positions about anti-

---

<sup>197</sup> Chemero, 'Radical embodied cognitive science', 145.

<sup>198</sup> Chemero, 'Radical embodied cognitive science', 145.

<sup>199</sup> Chemero, 'Radical embodied cognitive science', 149.

<sup>200</sup> Chemero, 'Radical embodied cognitive science', 145.

representationalism, but their metaphysics are utterly incompatible. So, we must press on.

#### 6.3.4. Von Uexküll

Von Uexküll's position is slightly hard to characterise, since he is tacit on both debates: representationalism and realism. Nonetheless, he makes metaphysical assumptions which are important to unpack. Von Uexküll appears to entertain a type of representational relativism – a type of idealism – where every organism constructs its own world and inhabits its own world. For von Uexküll, each organism inhabits its own 'soap bubble'.<sup>201</sup> This is the idealist's appeal to metaphysical foundations, which says that the image came before the world.<sup>202</sup>

However, intriguingly, von Uexküll's stance on metaphysics is not always clear. In one particular phrase, he appears to imply a realist position:

...out of the hundreds of stimuli radiating from the qualities of the mammal's body, only three become the bearers of receptor cues for the tick... Of all the influences that emanate from the mammal's body, only three become stimuli.<sup>203</sup>

Here, he appears to treat the hundreds of stimuli in a *mind-independent* way. Only three stimuli are perceptible to the tick, yet the others are held to persist nonetheless. This quote does not align with the view that the image came before the world. However,

---

<sup>201</sup> Von Uexküll, 'A stroll through the worlds of animals and men', 5.

<sup>202</sup> Varela, Thompson and Rosch, *The embodied mind*, 172.

<sup>203</sup> Von Uexküll, 'A stroll through the worlds of animals and men', 11.

elsewhere, von Uexküll's treatment of *Umwelten* and 'soap bubbles' is closely aligned with idealism. He is strongly pro-representationalist, since every organism represents a different perceptual world, and strongly anti-realist, since every organism is trapped in its own representational world.

One of Von Uexküll's most overt statements reads: 'There is no space independent of subjects'. There is no 'all-encompassing universal space'.<sup>204</sup> This can certainly be read in idealist terms, that the image comes before the world. Moreover, for von Uexküll, there is no world outside the images – no world outside the soap bubbles. So, it makes sense to place him squarely, more or less, in the idealist corner of the grid.

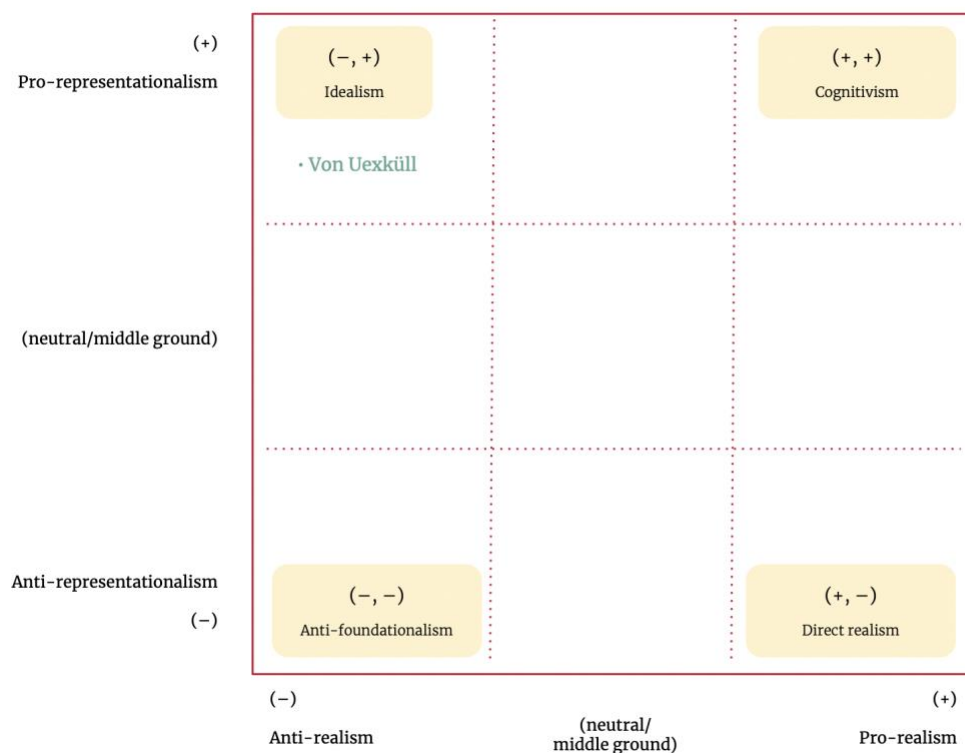


Figure 13: Von Uexküll is (tacitly) *pro-representationalist* and (tacitly) *anti-realist*

<sup>204</sup> Von Uexküll, 'A stroll through the worlds of animals and men', 29.

### 6.3.5. Hutto and Myin

Lastly, I will present the metaphysical agnostics: Hutto and Myin (§6.3.5) and Clark (§6.3.6). These authors argue that enactivism entails *neither* realism *nor* anti-realism. In their book, *Radicalizing enactivism*, Hutto and Myin present their radical enactivist position. They reject more ‘conservative’ forms of enactivism – O’Regan and Noë’s, and Thompson’s – which accept representations to a greater or lesser degree. By contrast, Hutto and Myin’s radical enactivism is truly opposed to cognitivism, since the two positions ‘logically exclude one another’.<sup>205</sup> They disagree, fundamentally, about the basic nature of mentality; cognitivism argues that it is *representational* and *contentful*, whereas radical enactivism argues that it is *non-representational* and *content-free*. Cognitivism says: ‘we must think in order to act’. Radical enactivism says: ‘we act before we think’.<sup>206</sup> Thus, Hutto and Myin’s enactivism completely rejects representations as a component of cognition.

Hutto and Myin’s metaphysical stance is particularly interesting. When it comes to metaphysics, Hutto and Myin take issue with the ‘more extravagant claims’ made by other enactivists.<sup>207</sup> They are referring to the anti-realist sentiment (from Varela, Thompson and Rosch) that there are multiple perceiver-dependent worlds, since each organism enacts their own world. Varela, Thompson and Rosch write that ‘our perceived world is constituted through complex and delicate patterns of

---

<sup>205</sup> Hutto and Myin, *Radicalizing enactivism*, 12.

<sup>206</sup> Hutto and Myin, *Radicalizing enactivism*, 12.

<sup>207</sup> Hutto and Myin, *Radicalizing enactivism*, 5.

sensorimotor activity'.<sup>208</sup> For Varela and his co-authors, this leads to the anti-realist conclusion that organisms enact and perceive *different* worlds – that the world does not exist independently of the perceiver. Hutto and Myin argue that this anti-realist conclusion does *not* follow from the enactivist programme. In fact, for Hutto and Myin, enactivism entails neither realism nor anti-realism.

Let us recall the disagreement of Gibson versus Varela, Thompson and Rosch. Is enactivism connected to anti-realism, since organisms participate in enacting the world (following Varela, Thompson and Rosch)? Or, is enactivism connected to realism, since organisms directly perceive their environment without mediating representations (following Gibson)? Interestingly, Hutto and Myin are steadfastly agnostic: 'In this book we remain neutral'.<sup>209</sup> They argue that the anti-realist verdict is *not* connected to their own enactivist arguments. Presumably, this means that the opposite is also true – that there is no clear inference in favour of realism.

---

<sup>208</sup> Varela, Thompson and Rosch, *The embodied mind*, 164.

<sup>209</sup> Hutto and Myin, *Radicalizing enactivism*, 5.

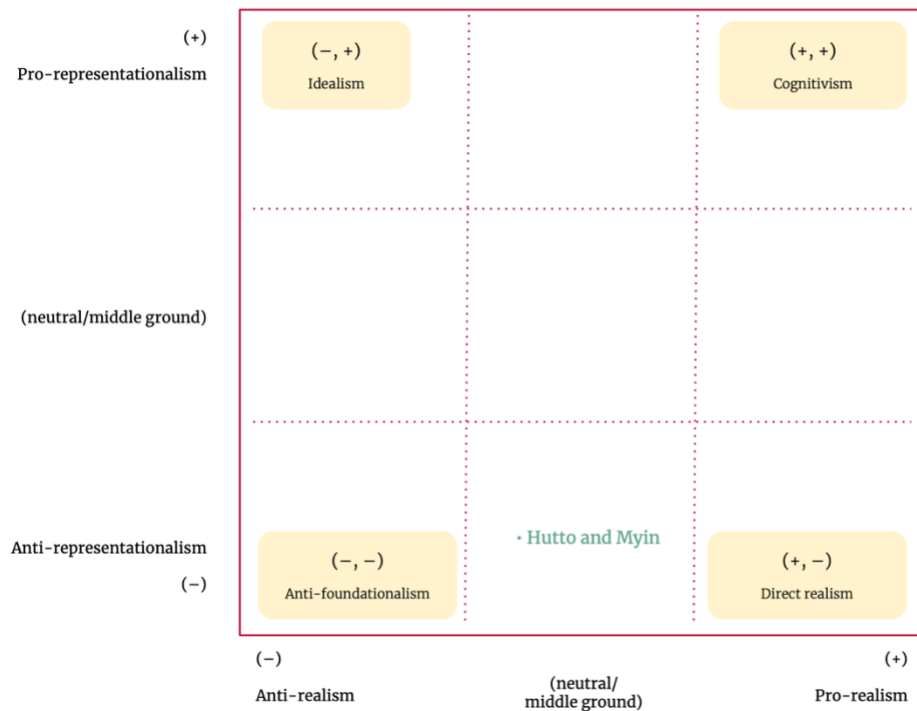


Figure 14: Hutto and Myin are *anti-representationalist* and *agnostic* about realism

Hutto and Myin devote less than half a page to dismissing this whole thread. It is certainly striking that they consider it to be so unimportant but, without further elaboration, we have little to respond to. I would welcome a longer entry from Hutto and Myin about the basis for their agnosticism; it is unsatisfying that they don't provide a more detailed discussion of realism and anti-realism. In §7, I will analyse the salience of their agnosticism. I will argue that their radical enactivism is a kind of anti-foundationalism that *does* carry metaphysical implications, so agnosticism is ruled out.



### 6.3.6. Clark

Clark appears in a temperate region of the grid: right at the centre. As I discussed in §3, he takes a middle ground on representationalism. Clark is opposed to both anti-representationalism and maximal representationalism. Instead, he supports ‘minimal representationalism’, consisting of personalised, action-oriented models.<sup>210</sup> In other words, representations that are *egocentric*. Earlier, I endorsed Clark’s minimal representationalism as the one that best explains the egocentric aspect of perception.

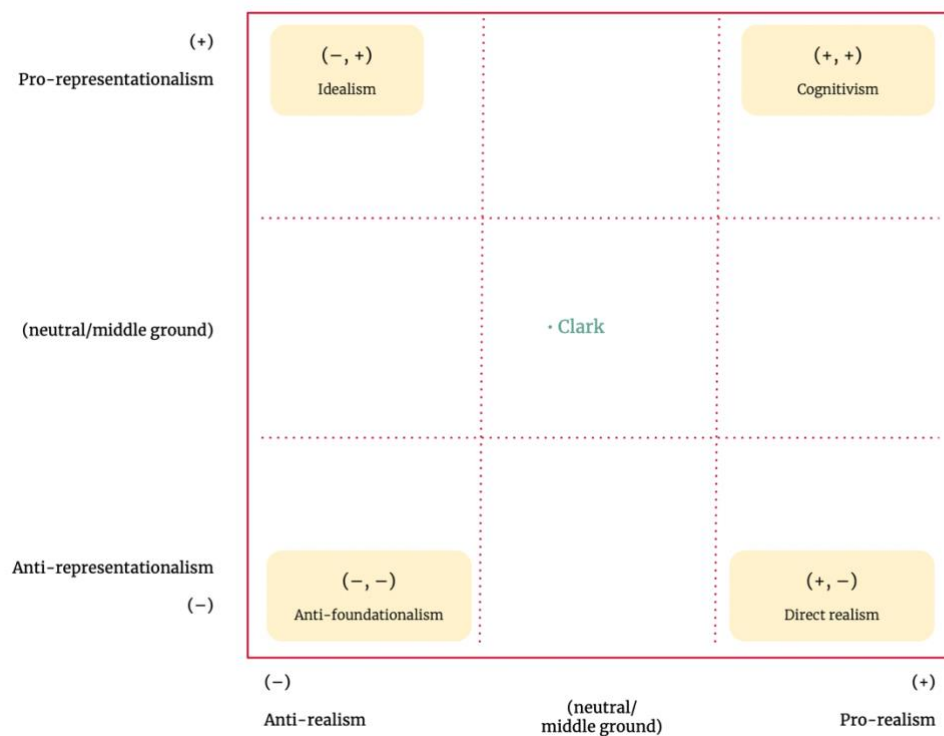


Figure 15: Clark takes a *middle ground* on representationalism and is *agonistic* about realism

<sup>210</sup> Clark, *Being there*, 174.

When it comes to realism, Clark – yet again – takes a middle ground. This is one of the three ways in which Clark distances himself from Varela, Thompson and Rosch. First, Varela, Thompson and Rosch completely reject the notion of representation, whereas Clark is ‘much more sympathetic’ to the role it can play in cognition. Second, Clark emphasises a different body of research – namely, real-world robotics and autonomous agent-theory – which might explain the divergence in both focus and conclusion. Lastly, Clark distances himself from the anti-realist ‘extension’ made by Varela and his co-authors.<sup>211</sup>

Clark does not want to make any claims about metaphysical reality; he does *not* object to ‘the idea that brains represent aspects of a real independent world’. Rather, he objects to the idea that brains construct *action-neutral* models of the world.<sup>212</sup> Clark is keenly focused on the nature of representations – whether they are maximal (action-neutral) or minimal (action-oriented). At the same time, he is agnostic on the topic of metaphysical realism. Clark’s agnosticism is reminiscent of Hutto and Myin’s agnosticism. They, too, distance themselves from the radical views of Varela and co. Hutto and Myin argue that anti-realism is *not* implied by enactivism – even by the most radical enactivism. Enactivism does not entail that the world is literally *constituted* and *brought into being* through sensorimotor activity. Enactivism does not require such anti-foundationalism about reality, they argue.

Chemero, in a review of Clark’s book, is not impressed with this sort of agnosticism. In his view, Clark sidesteps the anti-realism issue – a ‘casual sweeping

---

<sup>211</sup> Clark, *Being there*, 173.

<sup>212</sup> Clark, *Being there*, 173.

under the rug' – which lets the reader down. Chemero argues that anti-realism *does* follow from Clark's view of the mind. Even under minimal representationalism, our model of the world will be *action-oriented*. Thus, we do not perceive an objective, *action-neutral* world. For Chemero, *any* action-oriented influence will mean that we are cut off from perceiving the action-neutral world – i.e., metaphysical reality.<sup>213</sup>

Chemero's objection is relevant but not successful. This thesis project is aimed at explaining the egocentric (action-oriented) element of perception, and how that element is *compatible* with allocentricity and realism. Chemero argues that anti-realism follows from the claim that representations are action-oriented (as opposed to action-neutral). Expecting action-neutral representations sets the bar too high. It is not exactly a straw man, but it is not the picture of realism that I am striving to defend. Instead, I *accept* that perception has an indelibly egocentric, action-oriented character. Perception, necessarily, involves both egocentric and allocentric elements. In §5, I argued that humans can perceive allocentric properties, and abstract away from facts about their perspective, even if their perceptions are necessarily coloured by egocentricity. In §7, I will show how egocentricity and action-oriented representations are compatible with realism.

In one way, I am sympathetic to the spirit of Clark's agnosticism. He is simply arguing that action-oriented models do not entail *anti*-realism. In this respect, he disagrees with Chemero, and with Varela, Thompson and Rosch. I endorse that claim, since egocentricity does *not* entail anti-realism. At the same time, Clark refrains from making any substantive arguments in favour of realism. That is where my sympathy

---

<sup>213</sup> Chemero, 'A stroll through the world of animats and humans'.

with Clark subsidizes. In this thesis project, I will go one step further. It is not only the case that egocentricity does not entail anti-realism. As it will argue in §7, it is also the case that egocentricity is perfectly *compatible* with realism – an insight that Clark omits to acknowledge.

Before moving to the substantive section (§7), let us briefly review the procedural task, which is now complete. The grid has been populated. *Figure 16* shows the 3x3 matrix with all relevant positions marked. Now, with this comprehensive overview in hand, I can critically engage with each of the arguments. Considering all the candidates, I will endeavour to delineate a cogent and comprehensive account of the metaphysics of perception.

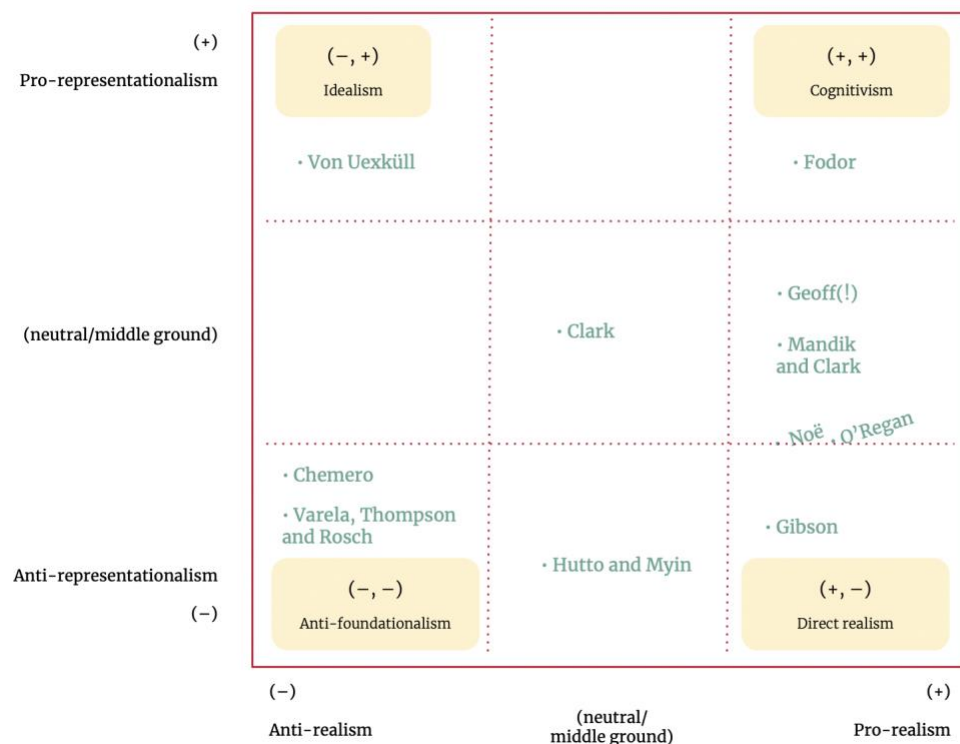


Figure 16: A 3x3 matrix comparing positions on representationalism with positions on realism

## 7. Enactivism and realism: Assessed

### 7.1. Assessing the realism debate

#### 7.1.1. In the spirit of symmetrical realism

In §5, I concluded that human perception is guided by both egocentricity and allocentricity. I referred to evidence from the neuroscience literature, such as the activity in the hippocampus and the increased grey matter in the hippocampus that accompany allocentric strategies for spatial navigation.<sup>214</sup> I also referred to empirical studies about object recognition. Morales and his team demonstrate that recognising the shape of a coin involves both egocentric and allocentric elements.<sup>215</sup>

Now, turning to the realism debate, I wish to frame the discussion in terms of egocentricity and allocentricity. This terminology gives us a unique tool for reframing the debate. In affirming metaphysical realism, I wish to do justice to *both* aspects of perception: the egocentric and the allocentric. My goal, in §7, is to account for the metaphysics of perceiver *and* world – to define a kind of symmetry between egocentricity and allocentricity. Building on the insights from §5, it is clear that a robust defence of realism must succeed in reconciling the two. I have already shown how enactivism astutely captures the dual aspect of perception, so an account of the metaphysics of enactivism must share this dual aspect.

In light of this dual aspect, it is illuminating to discuss Smith's position of 'symmetrical realism'. His symmetrical approach to metaphysics is on the right track.

---

<sup>214</sup> Colombo et al., 'Egocentric and allocentric spatial reference frames in aging', 606.

<sup>215</sup> Morales et al., 'Sustained representation of perspectival shape', 14873.

Smith wants to establish ‘a new metaphysics’. He wants to chart a middle way between naïve realism and ‘pure constructivism’.<sup>216</sup> His account encourages deference in both directions. On the one hand, following realism, we must have ‘epistemic deference to the world’. On the other hand, following constructivism, we must have deference to ‘human involvement in the world’.<sup>217</sup>

“We are here,” says the constructionist. “We are not the only things here,” says the realist.<sup>218</sup>

For Smith, these two insights – deference to human activity and deference to the world – are equally undeniable. Moreover, they are compatible with one another. There is no necessary conflict between the base claims of realism and constructivism.<sup>219</sup> I will reframe this in the language of ‘egocentricity’ and ‘allocentricity’. We can, at the same time, accept *both* egocentricity – the subject’s perspective – *and* allocentricity – the persistence and invariance of objects. We can have deference to *both* the fact of perspective *and* facts as they exist outside of perspectives.

I go no further in endorsing the content of Smith’s metaphysical account. I simply agree with the general spirit of symmetry in doing metaphysics. The dual nature of perception – egocentricity and allocentricity – must be met with dual deference – the perceiver and the world. Thus, I aim to account for the metaphysics of

---

<sup>216</sup> Smith, *On the origin of objects*, 3.

<sup>217</sup> Smith, *On the origin of objects*, 3.

<sup>218</sup> Smith, *On the origin of objects*, 88.

<sup>219</sup> Smith, *On the origin of objects*, 88.

perceiver and world, and to show that there is no contradiction between egocentricity and realism.

### 7.1.2. The egocentric profile

What do I mean by ‘the metaphysics of the perceiver’ or ‘the place of the perceiver in perception’? Going forwards, I will use the term *egocentric profile* to describe the inventory of facts about the perceiver’s body and location in the world – those facts that relate to the metaphysics of the perceiver. I will now outline the two core components of the egocentric profile: embodiment (facts about the physical body) and embeddedness (facts about vantage point). In developing my account of the egocentric profile, I was certainly inspired by the enactivist tradition, which emphasizes action, embodiment and embeddedness. Retaining these elements is part of my effort to rehabilitate enactivism.

Parthemore gives a nice summary of the enactivist conception of cognition, which says that cognition is ‘dynamic’ and ‘contingent’ rather than ‘fixed’. This contingency should be understood in terms of action: ‘*If I do this, then I will experience that*’. In this way, enactivism is closely related to Gibson’s ecological psychology.<sup>220</sup> This is a welcome reminder of the impact of the egocentric profile on perception. It has been the central argument of this thesis project that perception is *contingent* upon the egocentric profile: sensorimotor system, size, metabolism, vantage point, and so on. For the purposes of the realism debate, it is useful to subdivide this ‘profile’.

---

<sup>220</sup> Parthemore, ‘From a sensorimotor to a sensorimotor++ account of embodied conceptual cognition’, 143.

Enactivism and 4E cognition are linked to two core components: embeddedness (that cognition is shaped by the perceiver's physical location in the environment) and embodiment (that cognition is shaped by the perceiver's physical body).<sup>221</sup> Thus, these are two components in the egocentric profile: physical location (vantage point) and physical body (sensorimotor system, size, metabolism, and so on).

The account of sensorimotor contingencies is one thing that enactivism gets right. Enactivism gives an astute account of how perception is contingent upon facts about the perceiver, in terms of embodiment and embeddedness. Enactivism is a useful account precisely because it embraces talk of such egocentric elements. It embraces considerations about the metaphysics of the perceiver. As Parthemore puts it, 'the observer is always present and must be accounted for... sensorimotor contingencies... depend on the observer'.<sup>222</sup>

Throughout this thesis project, I have worked hard to understand and to engage with enactivism, because it is a useful way to *account for* the observer. As a result of this accounting, we are left with an inventory – a profile – of the observer's place in perception. The egocentric profile catalogues the contingencies which arise from the perceiver's embodiment and embeddedness. Thus, the egocentric profile informs our account of the metaphysics of perceiver, satisfying one half of the required symmetry. We must have deference to the place of the subject in the world.

Now, it remains to be shown that the egocentric profile is compatible with metaphysical realism. We must satisfy the other half of the symmetry – deference to

---

<sup>221</sup> Parthemore, 'From a sensorimotor to a sensorimotor++ account of embodied conceptual cognition', 146.

<sup>222</sup> *Ibid.*, 146–7.



the world. In the following subsections, I will argue for **Conclusion C2: Enactivism (and, in particular, the egocentric profile of a human perceiver) is compatible with metaphysical realism**. My argument for **Conclusion C2** is divided into two pillars: the argument from biology (§7.2) and the argument from geometry (§7.3). The two components of the egocentric profile – embodiment (§7.2) and embeddedness (§7.3) – require different metaphysical treatment. Thus, these two components inform the two pillars of my overall argument for metaphysical realism.

### **egocentricity + realism**

(pairwise interaction)

- Realists must account for both the place of the perceiver (egocentricity) and the persistence of the world (realism).
- A *symmetrical* account of the metaphysics of perception must reconcile the egocentric profile and metaphysical realism.

## **7.2. Embodiment: The argument from biology**

### **7.2.1. Egocentricity and anti-realism**

What can we say about biology and embodiment in connection to realism? In ‘Selective representing and world-making’, Mandik and Clark discuss ‘the thesis of

selective representing' and its implications for metaphysical realism.<sup>223</sup> According to the thesis of selective representing, representations of the world are perceiver-dependent; the world is represented differently by each organism, arising from differences in sensorimotor capacities and abilities. This is the view that each organism represents the world egocentrically, for its own specific purposes and needs. Thus, the thesis of selective representing describes our egocentric representations of the environment.

Mandik and Clark are discussing egocentricity, arising from embodiment. The sensorimotor system of each organism gives rise to different egocentric representations, based on different biological capacities and needs. As we saw in §6.3.4, this is von Uexküll's springboard for *anti*-realism. He writes: 'There is no space independent of subjects' – no 'all-encompassing universal space'.<sup>224</sup> For von Uexküll, facts about biology entail *anti*-realism. Every organism constructs a different perceptual world, and there is no reality outside of these multiple, perceiver-*dependent* worlds. This aligns closely with Khlentzos' definition of anti-realism, which I introduced in §6.1.2:

Anti-realists either doubt or deny the existence of the structure the metaphysical realist believes in or else doubt or deny its independence from our conceptions of it.<sup>225</sup>

---

<sup>223</sup> Mandik and Clark, 'Selective representing and world-making', 384.

<sup>224</sup> Von Uexküll, 'A stroll through the worlds of animals and men', 29.

<sup>225</sup> Khlentzos, 'Challenges to Metaphysical Realism'.

For von Uexküll, all that exists is many *egocentric* worlds. These worlds are not independent of the subjects who represent them. Each of these worlds is constrained by the sensorimotor system of the organism: the human, the gibbon, the goldfish, the tick, and so on. Thus, facts about biology (namely, differences in biology) entail metaphysical *anti*-realism.

Chemero makes a similar argument for anti-realism. If representations are egocentric ('action-oriented'), he argues, then humans have no 'special claim' that the world they inhabit is 'the *true reality* or *world-in-itself*'.<sup>226</sup> Facts about biology – and differences in biology between humans, gibbons, goldfish, etc. – give rise to multiple perceiver-dependent worlds. The human world, just like the world of any organism, is *not* 'an objective, action-neutral environment'.<sup>227</sup> Thus, for Chemero, facts about biology entails *anti*-realism about the objects of human perception. They are simply manifestations of our sensorimotor needs, and not features of an independent metaphysical reality. That is the case for egocentricity and anti-realism.

### 7.2.2. Egocentricity and realism

Mandik and Clark accept the egocentric nature of representations, arising from biology and the sensorimotor system. They rely on von Uexküll's tick as a pivotal example, while also mentioning the goldfish, the gibbon and the human. They, too, compare the *Umwelt* of the human and the *Umwelt* of the tick to reach an insight about egocentricity in perception. Each of these organisms – human, tick, goldfish or gibbon

---

<sup>226</sup> Chemero, 'A stroll through the world of animats and humans'.

<sup>227</sup> Chemero, 'A stroll through the world of animats and humans'.

– represents a different world, populated with a different set of ‘relevant’ stimuli. So, the effective environment of each organism (or species) is vastly different. The lifeworld of the tick is vastly different to the lifeworld of the human or the gibbon or the goldfish. Each organism represents differently and *selectively*, which gives rise to these differences in lifeworlds. In this way, they agree with von Uexküll and Chemero about the egocentric nature of representations.

However, Mandik and Clark argue forcefully that egocentricity does not entail metaphysical anti-realism:

Nonetheless, one must avoid the view that makes idealism rest on the following tautology: the only world that we represent is a world that is represented by us. Now, of course, the world represented by us is representation dependent in at least this sense: it depends on being represented by us for its being represented by us. But this can’t be what the realist and anti-realist are disagreeing about. We mention this point... only to point out the dangerous proximity between the tautologous version of anti-realism and the rhetoric employed in these discussions.<sup>228</sup>

Mandik and Clark are objecting to a ‘tautologous’ anti-realism employed by von Uexküll. Just because the tick’s representational world is dependent on the tick, it does not mean that the *metaphysical* world is dependent on the tick. For Mandik and Clark, there can still be a shared world outside of narrow representation, underpinning those representations. The crux of their paper aims to show that egocentricity (or, the thesis of selective representing) is fully compatible with realism.

---

<sup>228</sup> Mandik and Clark, ‘Selective representing and world-making’, 387.

Khlemtzos, in his entry 'Challenges to metaphysical realism', discusses egocentricity (in a way reminiscent of Mandik and Clark) and its implications for metaphysical realism.<sup>229</sup> He refers to egocentricity in terms of 'conceptual relativity' or 'conceptual pluralism' – the idea that there is a plurality of ways of representing, specific to different species. What does this mean for the existence of common-sense objects like 'trees, rocks and microbes'?<sup>230</sup> A sensible response to Khlemtzos' question should distinguish *entities themselves* from *representations of those entities*. We should recognise that representations of trees and rocks are subject-relative, in so far as they are represented selectively by humans (and other species). However, crucially, it is the *representations*, not the *entities*, which are subject-relative. Human representations of rocks have no place in the lifeworld of the tick, so they are mind-dependent, but the entities themselves – the rocks – are part of one, shared, mind-independent reality. Features of this mind-independent reality persist independently of the narrow perspective of any one species. Certain aspects, like concentrations of butyric acid, are invisible to us, whereas other aspects, like rocks and trees, are invisible to the tick.

Our vastly different representations of reality may be selective, but they are not mutually contradictory, as Mandik and Clark remind us. The narrow lifeworlds of different species are necessarily informed by a shared reality, which underpins those differences. Mandik and Clark make the case that egocentricity entails *realism*. They argue that egocentricity, arising from facts about biology, would not be possible without a prior, mind-independent world. The features which give rise to the

---

<sup>229</sup> Khlemtzos, 'Challenges to Metaphysical Realism'.

<sup>230</sup> Khlemtzos, 'Challenges to Metaphysical Realism'.

egocentric *Umwelt* of the tick are stimuli such as butyric acid, skin and heat – but these are real, objective, mind-independent constraints. It is *objective* differences between the sensorimotor systems of different species that give rise to differences in lifeworlds, and that give rise to selective representing. Without a shared, independent reality, there would be no objective differences, and selective representing would be impossible. As Mandik and Clark write, ‘some features may be subjective, but whether they are subjective is itself an objective matter’.<sup>231</sup> This is the *objective fact of subjectivity* or the *objective fact of egocentricity*.

As an example, let’s say that Bethany’s favourite colour is orange. Her preference for orange is *subjective*, because she subjectively believes that it is a beautiful colour. However, there is an *objective* fact as to her preference. Objectively, Bethany has a preference for the colour orange and, objectively, she has a belief about it being beautiful. That is the objective fact of subjectivity.

Let’s link Mandik and Clark’s argument back to enactivism and embodiment. According to enactivism, perception is constrained by sensorimotor contingencies – facts about the sensorimotor system of the organism (in relation to the environment). These facts relate to the physical body of the organism: size, metabolism, sensory apparatus, nervous system, and so on. Perception is inherently egocentric, because of these facts about biology. Mandik and Clark argue that these facts *entail* realism. A mind-independent world, which gives rise to neurobiology, is a necessary precondition to perceiving anything at all. The biological constraints which make perception different for humans versus ticks are *themselves* persistent. There is nothing

---

<sup>231</sup> Mandik and Clark, ‘Selective Representing and World-Making’, 391.

egocentric about the *conditions* that give rise to egocentricity. On the contrary, these conditions are symptomatic of a shared, mind-independent world.

Here, Mandik and Clark are making a very strong claim: that egocentricity entails realism. This argument comes with its flaws. In §7.2.3, I will consider an objection to their argument: that they are begging the question of realism. Ultimately, I will settle on a less strong claim about realism. It is simply the case that egocentricity is *compatible with* egocentricity. As I will outline, we do not have the basis to go as far as Mandik and Clark in saying that egocentricity *entails* realism.

### **7.2.3. Circularity: Begging the question of realism**

There is a clear objection to Mandik and Clark's argument. They are trying to *prove* metaphysical realism, and yet their argument seems to *presuppose* metaphysical realism, as a premise. This might amount to a vicious circularity. Interestingly, Gibson's argument for direct realism is symptomatic of the same circularity. Gibson uses the theory of direct perception as a *premise* in favour of direct realism, yet direct perception already *presupposes* direct realism. Thus, the following discussion of Mandik and Clark also applies to Gibson's account. Here is a clear statement of Mandik and Clark's argument:

Biological differences (in terms of neurobiology and sensorimotor systems) give rise to egocentricity, therefore the world (which gives rise to neurobiology and sensorimotor systems) must exist *prior to or independently of* egocentricity.

The very premise of this argument – biological differences give rise to egocentricity – is phrased in *realist* terms. Of course, the conclusion will lead to realism, since the argument has already been formulated in terms of realism. It is not likely that an anti-realist will be convinced by this argument. Anti-realists already disagree with the premise, since they contest the existence of mind-independent constraints (like neurobiology). For the anti-realist – like von Uexküll or Chemero – egocentricity is such that metaphysical structures (like neurobiology) may well be *mind-dependent*.

This objection to Mandik and Clark's argument is somewhat deflating. Their argument is a neat statement of realism, satisfying and agreeable to the ears of the realist. Yet, because it already begs the question, it does little to combat anti-realism. We need a different strategy. We must simply dilute the argument slightly, to make a more modest claim about realism. Using the same conceptual tools, we can make a more sure-footed argument: that metaphysical realism is *compatible with* egocentric facts about embodiment.

The argument with realism is compatible with facts about embodiment does not fall victim to the same circularity, since it only aims to explain embodiment in realist terms. We are setting out to explain embodiment in terms of realism, therefore we are allowed to use realism as a premise. This involves the more sober aim of showing that the egocentric profile is not at odds with metaphysical realism. Indeed, when it comes to embodiment, we can give a full explanation of the constraints that give rise to egocentricity, in terms of size, metabolism, neurobiology, sensorimotor system, and so on. Realism can be reconciled with egocentric facts about embodiment.



This is the first pillar of my overall argument for realism, **The argument from biology: Metaphysical realism is compatible with egocentric facts about embodiment.**

#### **7.2.4. Biological egocentricity: From the species to the individual**

I wish to consider one more objection to the argument from biology. I have argued that facts about biology are ‘egocentric’, relating to each individual organism. Yet, is this conception overly narrow? Don’t facts about biology relate to the level of the *species*, rather than the level of the *individual*? An objection might say that we should talk about ‘species-centricity’ rather than ‘egocentricity’, in the context of biological differences.

In response, I would argue that there are sufficient differences, at the level of the individual, which give rise to egocentric differences in perception. I would argue that each individual has a (slightly) different egocentric profile. For example, consider the example of Bethany and the bee, introduced in §5.3.2 and shown at *Figure 5*. There are differences between a human and a bee that give rise to differences in perceptual worlds. However, there are also egocentric facts specific to *Bethany* which constrain her egocentric profile, and make it different – at least slightly different – to other individual humans (in particular and on average). The length of Bethany’s legs is particularly relevant. She wants to cross the field to get to the ice-cream van, and she knows that it’s fourteen steps away. Bethany has mastered the (egocentric) sensorimotor contingencies connecting her footsteps and distances to ice-cream vans, and represents the distance in terms of her bodily abilities. This is already one fact specific to Bethany’s body, but not specific to human bodies in general.

Facts about Bethany's metabolism are also relevant. Bethany has always found delicious ice cream easy to metabolise; she would say it 'agrees with her'. On the other hand, Bethany's younger brother, Bill, is lactose intolerant. For him, ice-cream vans provide no promise or enticement. Sure, he heard the tinkling music of the ice-cream van approaching, but he chose to stay sitting in the grass at the far end of the field, reading his book. These biological differences between Bethany and Bill (foods they can metabolise) also constrain their egocentricity. The egocentric profile of Bethany arises from more than her more biology as a human being. Rather, there are egocentric facts specific to her eyes, legs, skin, metabolism, and so on. Thus, it is salient to talk about egocentric differences, arising from biology, at the level of the individual.

This discussion of the individual versus the species also raises a methodological question about the scope of this thesis project. Do my overall claims (realism and the egocentric profile) apply only to humans or to the wider animal kingdom? In brief, I am confining myself to claims about human perception. I do not have the evidential basis for making claims about other animals. At times, I have engaged with arguments – from Gibson, von Uexküll, and Mandik and Clark – that consider animal biology and animal perception. These are useful as guiding considerations, but I avoid making any conclusions about the ability of non-human animals to perceive the world. Interestingly, the evidence for allocentricity (associated with activity in the hippocampus) is not limited to humans but extends to rats.<sup>232</sup> This suggests that some of my core claims about enactivism and metaphysics might extend,

---

<sup>232</sup> Mandik, 'The neural accomplishment of objectivity'.

in some way, to other mammals. However, I do not have sufficient evidence to explore this possibility. In my overall conclusions, I deal only with human perception.

### **7.3. Embeddedness: The argument from geometry**

#### **7.3.1. Embeddedness and vantage points**

In §5, I explored the notion of ‘perspective’ at length. Following Noë, and following the evidence from Morales and co-authors, I concluded that perception has a ‘dual nature’, both egocentric and allocentric. I used the example of a plate. When Tiger Woods looks across the table at his friend’s empty dinner plate, he can perceive *both* the object-centred shape of the plate (circular) *and* the self-centred shape of the plate (elliptical). The plate really looks circular, and it really looks elliptical from his vantage point. The most important implication, with respect to the egocentric profile, is that the egocentric aspect of perception persists. For Tiger, the plate looks elliptical, because of his egocentric vantage point. Differences in egocentric perspectives give rise to differences in perspectival properties (P-properties), like perspectival size (P-size) and perspectival shape (P-shape).

How should we account for egocentricity arising from perspective? Can it be reconciled with metaphysical realism? Noë makes a convincing case for P-properties being objective and mind-independent, and therefore perfectly in keeping with realism. Crucially, for Noë, there is nothing subjective about P-properties. The relation between size and P-size is governed by precise mathematical laws.<sup>233</sup> They are facts

---

<sup>233</sup> Noë, *Action in perception*, 83.

about objective relations in the world, so there is no descent into subjectivism. Appearances are not subjective; they are not 'relations between things and your mind'. Rather, appearances are objective; they are facts about relations between objects and facts about particular vantage points.<sup>234</sup>

Again, Noë displays his tacit commitment to realism, explaining egocentricity in objectivist terms. Noë argues that we can give an objective account of egocentricity. There is a *fact* about a vantage point that you occupy in a physical environment. We can express these egocentric relations in terms of geometry. Tiger Woods, seated at the dinner party, occupies a physical location in the environment. The egocentric aspect of his perceptions – for example, how elliptical his friend's dinner plate looks – is constrained by facts about his vantage point. These constraints (embeddedness and facts about geometry) form the second component of the egocentric profile.

Let's link this discussion back to enactivism. The perceiver is *embedded* in the environment, insofar as he occupies a physical vantage point in that environment. Enactivism succeeds in accounting for this component of egocentricity. By emphasising the role of the perceiver – located in the environment, interacting with the environment – enactivism displays appropriate deference to the egocentric profile. Moreover, enactivism shows that these facts about geometry are perfectly compatible with realism. That is the second pillar of my argument for realism, **The argument from geometry: Metaphysical realism is compatible with egocentric facts about embeddedness.** We can account for embeddedness in objective, geometrical terms.

---

<sup>234</sup> Noë, *Action in perception*, 164.

Noë makes a broader point about the compatibility between egocentricity and realism. Noë argues that the ‘perceptual world’ is not a separate place from the ‘physical world’. It is simply the world from some standpoint.<sup>235</sup> Crucially, this means that ‘the perceptual world is not a subjective world’. Granted, different organisms may have limited access to the world, or to aspects of the world. And, granted, differences in perceptual systems determine the *availability* or *unavailability* of phenomena (arising from facts about biology). Yet we don’t bring the world into *existence*.<sup>236</sup> Noë argues for the compatibility of egocentricity and realism: ‘different animals inhabit different perceptual worlds, even though they inhabit the same physical world’.<sup>237</sup> This argument is pleasantly reminiscent of Khlentzos and of Mandik and Clark, all of whom argue that multiple representational schemes do not undermine metaphysical realism.

I have been complaining, in §6.3.2, that Noë’s commitment to realism is tacit and unsubstantiated. So, it might be surprising to discover that Noë has provided some solid tools for a realist account of perception. Indeed, especially for the account of perspectival properties as objective, we can look to Noë for support. Yet, Noë’s realism is still tacit – in that he does not explicitly discuss metaphysics, and does not provide any arguments to assuage sceptics and anti-realists. His account *relies on* realism, and *aligns with* realism, but does not explicitly defend realism. Therefore, it is necessary to spell out the explicit arguments in favour of realism. The fact that we can still draw on Noë’s work, in doing so, is more evidence that the metaphysics of

---

<sup>235</sup> Noë, *Action in perception*, 155.

<sup>236</sup> Noë, *Action in perception*, 156.

<sup>237</sup> Noë, *Action in perception*, 156.

enactivism is compatible with realism. That package provides a robust and salient account of human perception.

### 7.3.2. Embeddedness: Externalism and extensive cognition

I have just presented the second pillar of my argument for realism, **The argument from geometry: Metaphysical realism is compatible with egocentric facts about embeddedness.** Now, I wish to consider an objection to this argument. There is a further sense of ‘embeddedness’, which views the perceiver’s place in the world in terms of ‘the coupled, and even creative, dynamics of an organism-environment system’.<sup>238</sup> On this view, the perceiver is so *enmeshed* in the world that the dynamics of the perceiver and the dynamics of the world are not distinguishable. How does this sense of embeddedness fit into the realist’s picture? In this objection, I draw on Hutto and Myin, and Varela, Thompson and Rosch. Their more radical forms of enactivism entail an *externalist* view of cognition.

In §3.3, I introduced Hutto and Myin’s witty formulation for the dynamics of mentality. They argue that cognition arises from dynamic interactions which are ‘loopy’.<sup>239</sup> This rejects the received view that cognition involves linear relations between inputs and outputs, or between inner (mentality) and outer (physicality). Instead – because cognition is radically *embedded* and *extensive* – there is no way to make such distinctions as input/output, inner/outer, or mental/physical. Cognition arises from interactions between all parts of the brain, body and environment. These

---

<sup>238</sup> Mandik and Clark, ‘Selective Representing and World-Making’, 389.

<sup>239</sup> Hutto and Myin, *Radicalizing enactivism*, 6.

interactions form cascading feedback loops, where no part of the system has any prior or privileged role to play. It is certainly fitting to characterise these dynamics as 'loopy'.

Hutto and Myin are putting forward an *externalist* view of cognition. More precisely, their position is one of 'vehicle externalism': the claim that 'the subject's experience supervenes upon *more* than the events within the subject's body'.<sup>240</sup> According to the opposing view, *internalism*, the brain (and the body) has a privileged role. Mentality and cognition still originate in the brain, even if some processes can (sometimes) be extended. This is the internalism versus externalism debate. Internalists argue that the brain and body is a sufficient subvening base for cognition. Externalists argue that cognition supervenes upon the *environment* as well as the brain and body.

Hutto and Myin argue for a kind of radical externalism: that cognition is not merely extended, but fundamentally extensive. They write: 'minds are already, in their basic nature, *extensive* and wide-ranging'.<sup>241</sup> For Hutto and Myin, mentality – in all cases – arises from 'the extensive ways in which organisms interact with their environments'; mentality never refers to the activity of the brain in isolation.<sup>242</sup> Hutto and Myin's radical form of enactivism views the *interactions* between perceiver and environment as the subvening base for cognition. Mentality *never* arises from the brain in isolation; mentality *always* arises from the interaction between organism and

---

<sup>240</sup> Coates, *The metaphysics of perception*, 91.

<sup>241</sup> Hutto and Myin, *Radicalizing enactivism*, 7 (emphasis added).

<sup>242</sup> Hutto and Myin, *Radicalizing enactivism*, 7.

environment. This flies in the face of the traditional internalist position, which views the inner states of the subject (neural or otherwise) as sufficient for cognition.

Varela, Thompson and Rosch defend a similar radical externalism and embeddedness. As I discussed in §3.3, their metaphysical account involves a feedback loop between mind and world. Mind and world are enmeshed together, and the dynamics of one is not distinguishable from the dynamics of the other. They argue for ‘mutual specification’ – ‘mind and world specify each other’.<sup>243</sup> That is their much more radical conception of the embeddedness of the perceiver. Varela, Thompson and Rosch would also object to my argument from geometry. They would argue that embeddedness entails much more than mere vantage points. Embeddedness has metaphysical implications about the inseparability of perceiver and world. This carries anti-realist implications, as I will outline next.

### 7.3.3. Embeddedness and anti-realism

What is the relation between realism and maximally loopy cognition? A radically embedded account of cognition (like that of Varela, Thompson and Rosch or Hutto and Myin) entails anti-realism about the mind-independent world. More precisely, it entails onto-epistemological anti-foundationalism – the specific brand associated with Varela, Thompson and Rosch.<sup>244</sup>

For Varela, Thompson and Rosch, the dynamics of mind and world are intrinsically combined, so neither can be said to be *independent* of the other, and neither

---

<sup>243</sup> Varela, Thompson and Rosch, *The embodied mind*, 172.

<sup>244</sup> Vörös and Riegler, ‘A plea for not watering down the unseemly’, 3.



can be said to be more *fundamental* to the other. The dynamics of mind and dynamics of world are indistinguishable. There are no mind-independent objects and there are no object-independent minds.<sup>245</sup> This aligns with Khlentzos' definition of anti-realism:

Anti-realists either doubt or deny the existence of the structure the metaphysical realist believes in or else doubt or deny its independence from our conceptions of it.<sup>246</sup>

Varela, Thompson and Rosch's position is anti-realist because it holds that the world is *not* independent of minds. A radically embedded view of cognition entails such anti-realism. Alarm bells should now be ringing! *If* radically embedded cognition entails anti-realism, *then* we must find a strategy for refuting radically embedded cognition – if we are going to affirm realism. That is my next task, in §7.3.4.

It is time, at last, to refute Hutto and Myin's agnosticism. Hutto and Myin take a similar position to Varela, Thompson and Rosch about radically embedded cognition. Yet, they argue that this position carries no metaphysical implications either way.<sup>247</sup> That is clearly mistaken, since their type of radical ('loopy') enactivism holds that the dynamics of mind and dynamics of world are indistinguishable. If mind and world are enmeshed and indistinguishable, this rules out realism. Realism is the thesis that the world is *independent of* (and, therefore, *separate to*) the mind:

---

<sup>245</sup> Varela, Thompson and Rosch, *The embodied mind*, 172.

<sup>246</sup> Khlentzos, 'Challenges to Metaphysical Realism'.

<sup>247</sup> Hutto and Myin, *Radicalizing enactivism*, 5.

Metaphysical realism is the thesis that the objects, properties and relations the world contains, collectively: the structure of the world ... exists independently of our thoughts about it or our perceptions of it.<sup>248</sup>

Agnosticism is simply not available to Hutto and Myin. Radically embedded cognition entails metaphysical anti-realism. The thesis of realism requires the independence of the world. This reminds us of the important task ahead. In order to affirm realism, we must refute radically embedded cognition. We must opt for a more sober and conservative account of embeddedness – one which is compatible with realism.

#### **7.3.4. Embeddedness and realism**

Now, I must refute radical embeddedness. Luckily, I laid the foundations in §3, in my treatment of representationalism. Radically embeddedness (as put forward by Varela, Thompson and Rosch, and by Hutto and Myin) also entails anti-representationalism. Hutto and Myin explicitly reject any inner/outer distinction:

...there is no way to isolate properly mentality-constituting “inner” organismic responses from “outer” ones that allegedly stand over and against the former as mere causal contributions from the environment. On this model, there is no prospect of making any such principled division.<sup>249</sup>

---

<sup>248</sup> Khlentzos, ‘Challenges to Metaphysical Realism’.

<sup>249</sup> Hutto and Myin, *Radicalizing enactivism*, 6.

Since the organism is radically embedded and enmeshed in the environment, there is no distinction between the 'inner' states of the organism versus the 'outer' states of the environment. The same, of course is true for Varela, Thompson and Rosch's account. Their account is built around the fundamental inseparability and 'structural coupling' of perceiver and world:

...this fundamental circularity... we found that we could discern no subjective ground, no permanent and abiding ego-self... we found a world enacted by our history of structural coupling.<sup>250</sup>

Thus, in order to refute radical embeddedness, we must show that mind and world *can* be distinguished. We must show that mind and world are *not* intrinsically coupled. It will be enough to show that perceivers have 'inner' representational states, which are distinct from 'outer' environmental entities. Thus, my endorsement of minimal representationalism (§3.5) is crucial for demonstrating the existence of distinctly 'inner' states. Here, I will re-affirm my case for the existence of (some) (minimal) representations, as a way to refute radical embeddedness.

Clark, the proponent of minimal representationalism, also discusses the notion of embeddedness. He considers an objection to representationalism, from radical embeddedness: the case of 'continuous reciprocal causation'. In this case, the organism is not insulated from its environment. Rather, it is embedded in its

---

<sup>250</sup> Varela, Thompson and Rosch, *The embodied mind*, 217.

environment – merged through ‘continuous, mutually modulatory influences, linking brain, body and world’.<sup>251</sup> This characterisation is reminiscent of what Hutto and Myin called ‘loopy’ dynamics.<sup>252</sup> Loopy dynamics (‘dense reciprocal causal influence’) mean that we cannot distinguish between inputs and outputs. There is no sharp divide between mind and world, and there is no linear input-output relation between the two. Instead, mind and world are enmeshed together, and they are co-constituted through feedback loops.<sup>253</sup>

Clark accepts that this extreme case is possible. In certain situations, he admits, mind and world may be so densely coupled that there are *no* identifiable inner states, of which to speak. For Clark, these cases are interesting, but do *not* constitute a serious challenge to representationalism. Crucially, they are exceptional cases, not any relevant majority.<sup>254</sup> So, while radically embeddedness is sometimes possible, it is merely an exception, and not a defining feature of cognition. Often – even with *some* loopy dynamics at play – there are also *some* essentially inner states. In these intermediate cases, it is still coherent and constructive to treat such inner states as representations.<sup>255</sup>

At the opposite extreme, the ‘representation-hungry’ cases provide a stronger defence for representationalism. When an organism (a human) must reason about the abstract or the non-existent, we must unavoidably refer to the *internal* dynamics of

---

<sup>251</sup> Clark, *Being there*, 163.

<sup>252</sup> Hutto and Myin, *Radicalizing enactivism*, 33.

<sup>253</sup> Clark, *Being there*, 165

<sup>254</sup> Clark, *Being there*, 166; 174–5.

<sup>255</sup> Clark, *Being there*, 164–5.

that cognitive agent.<sup>256</sup> These ‘representation-hungry’ problems require a distinction between internal and external dynamics, and it is coherent to make such a distinction.

Clark’s argument provides further support for minimal representations – inner states which are distinct from outer events – as I argued at length in §3. The internal dynamics of a cognitive agent are a vehicle for representing outer states of the world, especially when reasoning about states of affairs that are: abstract, non-existent, past, future or spatially distant. Thus, we are in a position to refute radical embeddedness, as expressed by Hutto and Myin and by Varela, Thompson and Rosch. Contrary to Hutto and Myin, contrary to Varela, Thompson and Rosch, the perceiver is not enmeshed in the environment to the point that inner/outer dynamics are indistinguishable. Rather, there are distinctly inner states (minimal representations) which carry information about outer events in the environment. On my more moderate conception of embeddedness, perception is still constrained by features of the physical environment, but those features are distinguishable from the mental states of the organism.

Accordingly, we arrive at a metaphysical picture that is compatible with realism. Since inner states are distinct from outer states, the world can be held to be *independent of minds*. My account of embeddedness rests on the argument for minimal representations, put forward in §3 and §4, and this is an account of embeddedness that is compatible with realism. Crucially, this adds support for the second pillar of my argument for realism, **The argument from geometry: Metaphysical realism is compatible with egocentric facts about embeddedness.** I have shown that radical

---

<sup>256</sup> Clark, *Being there*, 175.

embeddedness can be refuted and, moreover, I have provided an account of embeddedness that is compatible with metaphysical realism.

#### 7.4. Closing comments on metaphysical realism

In §6 and §7, I set out to investigate **Sub-question C: *What is the link between enactivism and realism, as claimed by other authors and as justified by the evidence?***

I entered the realism debate, employing the terminology of egocentricity and allocentricity. I brought a unique *perspective* to this debate, with a particular focus on the metaphysics of enactivism. In §6, I mapped out the interrelations between the realism and representationalism debates, and presented a visual matrix showing the positions of a wide range of philosophers (*Figure 16*). I argued for **Conclusion C1: Enactivists disagree about metaphysical implications, claiming realism, anti-realism or agnosticism.** This disagreement highlighted a need to analyse the metaphysics of enactivism. In §7, I engaged critically with a range of arguments for and against realism. Ultimately, I argued that it is possible to give an account of enactivism (and perception, more broadly) that is compatible with metaphysical realism.

I approached my account of realism with the spirit of symmetry, giving equal deference to the ‘egocentric profile’ of the perception. According to my formulation, the egocentric profile is made up of two components: embodiment and embeddedness. Each of these components constrains perception and gives rise to egocentricity. Yet, I showed that the egocentric profile can be reconciled with metaphysical realism. First, I discussed embodiment, and outlined **The argument**

**from biology: Metaphysical realism is compatible with egocentric facts about embodiment.** This claim is sober and guarded, in light of an objection about begging the question of realism. Second, I discussed embeddedness, and made **The argument from geometry: Metaphysical realism is compatible with egocentric facts about embeddedness.** Here, I considered a challenging objection from Varela, Thompson and Rosch, and from Hutto and Myin, that purported to undermine the independence of the world from minds. In responding to this objection, I presented an account of embeddedness and minimal representationalism that is compatible with metaphysical realism. Thus, having established both pillars of my argument – the argument from biology and the argument from geometry – I reach the culmination of my account, **Conclusion C2: Enactivism (and, in particular, the egocentric profile of a human perceiver) is compatible with metaphysical realism.** That is my overall account of the *metaphysics of enactivism*.

## Overview of Block C: Enactivism and metaphysics

(§6, §7)

- **Sub-question C:** *What is the link between enactivism and realism, as claimed by other authors and as justified by the evidence?*
- **Conclusion C1:** Enactivists disagree about metaphysical implications, claiming realism, anti-realism or agnosticism.
- **Conclusion C2:** Enactivism (and, in particular, the egocentric profile of a human perceiver) is compatible with metaphysical realism.
  - **The argument from biology:** Metaphysical realism is compatible with egocentric facts about embodiment.
  - **The argument from geometry:** Metaphysical realism is compatible with egocentric facts about embeddedness.



## 8. General conclusion

### 8.1. The metaphysics of enactivism

#### 8.1.1. Enactivism

In this thesis project, I set out to investigate the metaphysics of enactivism. In §1, I introduced the following research question:

**The metaphysics of enactivism: Realism and the egocentric profile**

Research question

*In the context of **enactivism** and **the egocentric profile**,  
what are the implications (or what is the evidence) for metaphysical **realism**?*

I divided this research question into three sub-questions (A, B and C), and tackled the problem in three blocks (A, B and C).

First, I asked **Sub-question A: *To what extent has enactivism challenged or influenced representationalism?*** I introduced cognitivism, the ‘orthodoxy’, as espoused by Fodor and Pylyshyn. Then, I considered a range of enactive approaches to representation – from Varela, Thompson and Rosch, to O’Regan and Noë – all of which cite Gibson as a historical influence. I concluded that enactivists make a salient case for the role of action in perception. However, I also concluded that enactivists go too far in rejecting representations. The evidence supports minimal representations:

partial, personalised, action-oriented models of the world. The existence of minimal, action-oriented models is compatible with the role of action in perception. Thus, I reached **Conclusion A1: Perception is guided by both action and minimal representations.**

I pointed to a further flaw in O'Regan and Noë's account; it suffers from an internal inconsistency. On the one hand, they use anti-representationalist rhetoric. On the other hand, their account of sensorimotor contingencies requires perceivers to represent their expectations, and requires perceivers to represent an egocentric index. So, their account relies upon minimal representations. This inconsistency must be 'rehabilitated', since *on ne peut pas avoir le beurre et l'argent du beurre* (you can't have the butter and the money for the butter). Thus, I presented **Conclusion A2: Enactivism must be rehabilitated to include minimal representations.** In fact, after some gentle rehabilitation, an enactivist account that includes both action and minimal representations becomes a robust account of perception – one that is supported by the philosophical and empirical evidence.

### 8.1.2. Egocentricity

The most important insight from Block A was the importance of minimal representations in perception. So, in Block B, I set out to investigate the nature of these representations. What does it mean to say that a model is personal and action-oriented? I argued that a 'personal' model should be understood in terms of egocentricity. In this section, I asked **Sub-question B: *To what extent should human perception be understood as 'egocentric' or 'allocentric'?*** Here, my overall motivation

of symmetry is relevant. I wanted to investigate, and to account for, both aspects of perception: both egocentricity and allocentricity. An *asymmetrical* account poses a dangerous challenge. For example, von Uexküll's anti-realism flows from his view that perception is purely personal and egocentric, that every organism is trapped in their own representational world. In response to this challenge, I provided evidence for both egocentricity (associated with activity in the posterior parietal cortex) and allocentricity (associated with activity in the hippocampus). When it comes to spatial navigation, humans employ a mixture of egocentric and allocentric strategies for navigating the world. When it comes to object recognition, humans can detect both perspectival shape (egocentric) and distal shape (allocentric). Thus, the evidence supports **Conclusion B1: Perception is guided by both egocentric and allocentric elements.**

This thesis project is concerned with the metaphysics of enactivism, so I traced the implications of egocentricity and allocentricity for the enactivist account of perception. I uncovered strong links. O'Regan and Noë's notion of sensorimotor contingencies should be understood as egocentric representations, since the perceiver represents his own bodily abilities and potentials for action: *if I do this, then I will perceive that*. These contingencies are egocentric, since they are specific to the organism's embodiment and embeddedness. Moreover, when it comes to object recognition, we can explain Noë's account of the 'dual nature' of perception in terms of egocentricity and allocentricity. Humans can detect perspectival properties because perception is constrained by embodiment and embeddedness (egocentric), and humans can detect distal properties because we also possess sensorimotor knowledge

– an understanding of the invariant nature of relations and properties outside of our own perspective (allocentric). Thus, the key insights from enactivism can be stated in terms of egocentricity and allocentricity, adding support to **Conclusion B2: Enactivism should be understood in terms of egocentricity and allocentricity.**

### 8.1.3. Metaphysics

The link between enactivism and egocentricity is key for the final element in my thesis project: metaphysical realism. In Block C, I asked **Sub-question C: *What is the link between enactivism and realism, as claimed by other authors and as justified by the evidence?*** There is a surprising divergence between enactivists as to the metaphysics of enactivism: O'Regan and Noë are realists; Varela, Thompson and Rosch are anti-realists; Hutto and Myin, and Clark, are agnostics. I presented this insight in the form of **Conclusion C1: Enactivists disagree about metaphysical implications, claiming realism, anti-realism or agnosticism.** Such widespread disagreement called out for further investigation.

In the final section of this thesis project, I set out to account for the metaphysics of enactivism in realist terms – I set out to reach **Conclusion C2: Enactivism (and, in particular, the egocentric profile of a human perceiver) is compatible with metaphysical realism.** I argued that the egocentric profile is made up of two components: embodiment and embeddedness. Thus, my account was divided into two pillars, treating these two components separately. First, I made **The argument from biology: Metaphysical realism is compatible with egocentric facts about embodiment.** I rejected a stronger claim (from Mandik and Clark, and from Gibson)

that biological facts *entail* realism; this argument suffers from a vicious circularity. Instead, I simply set out to show that the egocentricity arising from embodiment is compatible with realism. This argument does not fall victim to the same circularity, since it explicitly operates in realist terms. On a realist view, there is nothing problematic about biological egocentricity, since differences between organisms (size, metabolism, sensorimotor system, and so on) arise from mind-independent constraints.

My second pillar was **The argument from geometry: Metaphysical realism is compatible with egocentric facts about embeddedness.** I began by explaining perspective and vantage point in realist terms. Then, I considered a more radical conception of embeddedness, whereby the organism is *enmeshed* with the world, meaning that the dynamics of the organism and the dynamics of the environment are fundamentally entangled and coupled. This radical embeddedness – espoused by Varela, Thompson and Rosch, and by Hutto and Myin – entails anti-realism, because the world is not *independent* of the organism. Independence is a key component in the definition of metaphysical realism. Contrary to Hutto and Myin’s professed agnosticism, their position does entail metaphysical anti-realism. In response, in order to defend realism, I had to show that the world *is* independent of the organism, that there *is* a distinction between the dynamics of the organism and the dynamics of the world. Returning to my argument for minimal representationalism from Block A, I argued that there must be ‘inner’ mental states which carry information about ‘outer’ states of affairs. Some type of inner/outer distinction is entailed by the notion of representation. Thus, I refuted radical embeddedness, and formulated a more

moderate conception of embeddedness, one which is compatible with metaphysical realism. Perception may be constrained by features of the organism's environment, but there is still a distinction between inner states versus the outer world.

#### **8.1.4. Realism and the egocentric profile**

Combining these blocks, I have now expounded an account of the metaphysics of enactivism. My account draws heavily on O'Regan, Noë and Clark, and yet it departs from each of these authors in significant ways. On the one hand, I endorsed Clark's minimal representationalism, which became a key part of my defence of metaphysical realism. So, on the other hand, I rejected Clark's agnosticism about metaphysics.

My overall account might be seen as an attempt to 'rehabilitate' O'Regan and Noë's sensorimotor enactivism. The first major flaw in their account – an inconsistent position about the role of representations – can be resolved by incorporating minimal representations. The other major flaw is that they do not provide an explicit defence of realism. By adding these elements, I showed that a rehabilitated enactivism (something close to what O'Regan and Noë intended) can provide a robust and empirically supported account of the metaphysics of perception. Most crucially, this rehabilitated enactivism is compatible with metaphysical realism. Thus, we can account for action, embodiment, embeddedness and how the perceiver relates to the mind-independent world. We can account for both the egocentric profile and metaphysical realism.

## 9. Bibliography

Allen, C. and Neal, J. 'Teleological notions in biology', in Zalta, E.N. (ed.) *The Stanford Encyclopedia of Philosophy* (Spring 2020) [online]. Available at: <https://plato.stanford.edu/archives/spr2020/entries/teleology-biology> (Accessed 1 June 2021).

Allen, G. (host), 'Pete Mandik: Perceiving the world', *Extrapolator*, podcast audio, forthcoming. Recorded on 2 July 2021.

Bishop, J.M. and Martin, A.O. 'Contemporary sensorimotor theory: A brief introduction', in Bishop, J.M. and Martin, A.O. (eds.) *Contemporary sensorimotor theory* (New York, NY: Springer, 2014), 1–22.

Chemero, A. 'A stroll through the worlds of animats and persons: Review of *Being there: Putting brain, body and world together again* by Andy Clark', *Psyche* [online], 4:14 (1998). Available at: <http://journalpsyche.org/files/0xaa4a.pdf> (Accessed 1 December 2020).

Chemero, A. 'Radical embodied cognitive science', *Review of General Psychology*, 17:2 (2013), 145–150.

Clark, A. *Being there: Putting brain, body and world together again* (Cambridge, MA: MIT Press, 1997).

Clark, A. and Chalmers, D. 'The extended mind', *Analysis*, 58:1 (1998), 7–19.

Coates, P. *The metaphysics of perception: Wilfrid Sellars, critical realism and the nature of experience* (New York, NY: Routledge, 2007).

Colombo, D. Serino, S., Tuena, C., Pedroli, E., Dakanalis, A., Cipresso, P. and Riva, G. 'Egocentric and allocentric spatial reference frames in aging: A systematic review', *Neuroscience and Biobehavioural Reviews*, 80 (2017), 605–621.

Flament, D., Ellermann, J.M., Kim, S.-G., Ugurbil, K. and Ebner, T.J. 'Functional magnetic resonance imaging of cerebellar activation during the learning of a visuomotor dissociation task', *Human Brain Mapping*, 4 (1996), 210–226.

Fodor, J. *LOT 2: The language of thought revisited* (Oxford University Press, 2008).

Fodor, J. and Pylyshyn, Z. 'How direct is visual perception? Some reflections on Gibson's "ecological approach"', in Noë, A. and Thompson, E. (eds.) *Vision and mind: Selected readings in the philosophy of perception* (Cambridge, MA: MIT Press, 2004), 167–227.

Gibson, J.J. *The ecological approach to visual perception* (London: Lawrence Erlbaum Associates, 1986).

Gibson, J.J. 'New reasons for realism', *Synthese*, 17 (1967), 162–172.

Gibson, J.J. 'A theory of direct visual perception', in Noë, A. and Thompson, E. (eds.) *Vision and mind: Selected readings in the philosophy of perception* (Cambridge, MA: MIT Press, 2004), 77–89.

Grush, R. 'Self, world and space: The meaning and mechanisms of ego- and allocentric spatial representation', *Brain and Mind*, 1 (2000), 59–92.

Hutto, D.D. and Myin, E. *Radicalizing enactivism: Basic minds without content* (Cambridge, MA: MIT Press, 2013).

Jahn, G., Knauff, M. and Johnson-Laird, P.N., 'Preferred mental models in reasoning about spatial relations', *Memory & Cognition*, 35 (2007), 2075–2087.

Khrentzos, D. 'Challenges to metaphysical realism', in Zalta, E.N. (ed.) *The Stanford Encyclopedia of Philosophy* (Winter 2016) [online]. Available at: <https://plato.stanford.edu/archives/win2016/entries/realism-sem-challenge> (Accessed 1 December 2020).

Mandik, P. 'The neural accomplishment of objectivity', in Poirier, P. and Faucher, L. (eds.) *Des neurosciences à la philosophie: Neurophilosophie et philosophie des neurosciences* (Paris: Éditions Syllepse, 2008). Available at: <http://www.petemandik.com/philosophy/papers/accompobj.html> (Accessed 1 April 2021).

Mandik, P. and Clark, A. 'Selective representing and world-making', *Minds and Machines*, 12:3 (2002), 383–395.



Meilinger, T. and Vosgerau, G. 'Putting egocentric and allocentric into perspective', in Hölscher, C., Shipley, T.F., Belardinelli, M.O., Bateman, J.A. and Newcombe, N.S. (eds.) *Spatial Cognition VII* (Heidelberg: Springer, 2010), 207–221.

Morales, J., Bax, A. and Firestone, C. 'Sustained representation of perspectival shape', *Proceedings of the National Academy of Sciences (PNAS)*, 117:26 (2020), 14873–14882.

Neurath, O. 'Anti-Spengler', in Neurath, M. and Cohen, R.S. (eds.) *Empiricism and sociology, Vienna Circle collection, vol. I* (Dordrecht: D. Reidel, 1973), 158–213.

Newen, A., Gallagher, S. and De Bruin, L. '4E cognition: Historical roots, key concepts, and central issues', in Newen, A., Gallagher, S. and De Bruin, L. (eds.) *The Oxford handbook of 4E cognition* (Oxford University Press, 2018), 1–16.

Noë, A. *Action in perception* (Cambridge, MA: MIT Press, 2004).

Noë, A. 'Concept pluralism, direct perception, and the fragility of presence', in Metzinger, T. and Windt, J.M. (eds.) *Open MIND*, 27(T) (Frankfurt am Main: MIND Group, 2015), 1–15.

Noë, A. 'Experience without the head', in Gendler, T.S. and Hawthorne, J. (eds.) *Perceptual experience* (Oxford University Press, 2006), 411–433.

Noë, A. 'What does change blindness teach us about consciousness?', *Trends in Cognitive Sciences*, 9:5 (2005), 218–218.

Noë, A. and O'Regan, K. 'On the brain-basis of visual consciousness: A sensorimotor account', in Noë, A. and Thompson, E. (eds.) *Vision and mind: Selected readings in the philosophy of perception* (Cambridge, MA: MIT Press, 2004), 567–598.

Noë, A. and Thompson, E. 'Introduction', in Noë, A. and Thompson, E. (eds.) *Vision and mind: Selected readings in the philosophy of perception* (Cambridge, MA: MIT Press, 2004), 1–14.

O'Regan, K. 'Solving the "real" mysteries of visual perception: The world as an outside memory', *Canadian Journal of Psychology*, 46 (1992), 461–488.

O'Regan, K. and Noë, A. 'A sensorimotor account of vision and visual consciousness', *Behaviour and Brain Sciences*, 24 (2001), 939–1031.

Parthemore, J. 'From a sensorimotor to a sensorimotor++ account of embodied conceptual cognition', in Bishop, J.M. and Martin, A.O. (eds.) *Contemporary sensorimotor theory* (New York, NY: Springer, 2014), 137–158.

Smith, C.B. *On the origin of objects* (Cambridge, MA: MIT Press, 1996).

Thompson, E. *Mind in life: Biology, phenomenology, and the sciences of mind* (Harvard University Press, 2007).

Turvey, M.T. 'Contrasting orientations to the theory of visual information processing', *Psychological Review*, 84 (1977), 67–88.

Varela, F.J., Thompson, E. and Rosch, E. *The embodied mind: Cognitive science and human experience* (Cambridge, MA: MIT Press, 2016).

Von Uexküll, J. 'A stroll through the worlds of animals and men: A picture book of invisible worlds', in Schiller, C.H. (ed. and trans.) *Instinctive behavior* (New York: International Universities Press, 1957), 5–80.

Vörös S. and Riegler A. 'A plea for not watering down the unseemly: Reconsidering Francisco Varela's contribution to science', *Constructivist Foundations*, 13:1 (2007), 1–10.

Wallis, C. and Wright, W. 'Enactivism's vision: Neurocognitive basis or neurocognitively baseless?', in Bickle, J. (ed.) *The Oxford handbook of philosophy and neuroscience* (New York, NY: Oxford University Press, 2009), 251–308.

Wiener, J.M., Kmecova, H. and de Condappa, O. 'Route repetition and route retracing: Effects of cognitive aging', *Frontiers in Aging Neuroscience*, 4 (2012), 1–7.